arXiv:1603.03304v4 [cs.CV] 25 Jan 2017

Template Matching via Densities on the Roto-Translation Group

Erik J. Bekkers, Marco Loog, Bart M. ter Haar Romeny, and Remco Duits

Abstract—We propose a template matching method for the detection of 2D image objects that are characterized by orientation patterns. Our method is based on data representations via orientation scores, which are functions on the space of positions and orientations, and which are obtained via a wavelet-type transform. This new representation allows us to detect orientation patterns in an intuitive and direct way, namely via cross-correlations. Additionally, we propose a generalized linear regression framework for the construction of suitable templates using smoothing splines. Here, it is important to recognize a curved geometry on the position-orientation domain, which we identify with the Lie group SE(2): the roto-translation group. Templates are then optimized in a B-spline basis, and smoothness is defined with respect to the curved geometry. We achieve state-of-the-art results on three different applications: detection of the optic nerve head in the retina (99.83% success rate on 1737 images), of the fovea in the retina (99.32% success rate on 1616 images), and of the pupil in regular camera images (95.86% on 1521 images). The high performance is due to inclusion of both intensity and orientation features with effective geometric priors in the template matching. Moreover, our method is fast due to a cross-correlation based matching approach.

Index Terms—template matching, multi-orientation, invertible orientation scores, optic nerve head, fovea, retina

1 Introduction

We propose a cross-correlation based template matching scheme for the detection of objects characterized by orientation patterns. As one of the most basic forms of template matching, cross-correlation is intuitive, easy to implement, and due to the existence of optimization schemes for real-time processing a popular method to consider in computer vision tasks [1]. However, as intensity values alone provide little context, cross-correlation for the detection of objects has its limitations. More advanced data representations may be used, e.g. via wavelet transforms or feature descriptors [2], [3], [4], [5]. However, then standard cross-correlation can usually no longer be used and one typically resorts to classifiers, which take the new representations as input feature vectors. While in these generic approaches the detection performance often increases with the choice of a more complex representation, so does the computation time. In contrast, in this paper we stay in the framework of template matching via cross-correlation while working with a contextual representation of the image. To this end, we lift an image $f : \mathbb{R}^2 \to \mathbb{R}$ to an *invertible orientation score* $U_f : \mathbb{R}^2 \rtimes S^1 \to \mathbb{C}$ via a wavelet-type transform using certain anisotropic filters [6], [7].

An orientation score is a complex valued function on the extended domain $\mathbb{R}^2 \rtimes S^1 \equiv SE(2)$ of positions and orientations, and provides a comprehensive decomposition of an image based on local orientations, see Fig. 1 and 2. Cross-correlation based template matching is then defined via $\mathbb{L}_2$ inner-products of a template $T \in \mathbb{L}_2(SE(2))$ and an orientation score $U_f \in \mathbb{L}_2(SE(2))$. In this paper, we learn templates T by means of generalized linear regression.

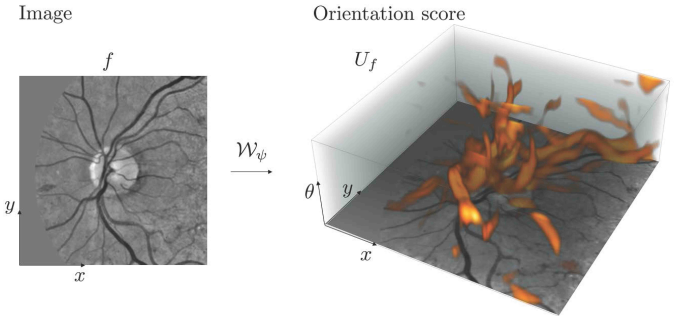

Fig. 1. A retinal image f of the optic nerve head and a volume rendering of the orientation score U_f (obtained via a wavelet transform $\mathcal{W}_\psi$).

In the $\mathbb{R}^2$-case (which we later extend to orientation scores, the $SE(2)$-case), we define templates $t \in \mathbb{L}_2(\mathbb{R}^2)$ via the optimization of energy functionals of the form

$$t^* = \underset{t \in \mathbb{L}_2(\mathbb{R}^2)}{\operatorname{argmin}} \{E(t) := S(t) + R(t)\}, \qquad (1)$$

where the energy functional $E(t)$ consists of a data term $S(t)$, and a regularization term $R(t)$. Since the templates optimized in this form are used in a linear cross-correlation based framework, we will use inner products in S, in which case (1) can be regarded as a generalized linear regression problem with a regularization term. For example, (1) becomes a regression problem generally known under the

- *E.J. Bekkers and B.M. ter Haar Romeny are with the department of Biomedical Engineering, Eindhoven University of Technology (TU/e), the Netherlands. E-mail: {e.j.bekkers,b.m.terhaarromeny}@tue.nl*
- *M. Loog is with the Pattern Recognition Laboratory, Delft University of Technology, the Netherlands. E-mail: m.loog@tudelft.nl*
- *B.M. ter Haar Romeny is also with the department of Biomedical and Information Engineering, Northeastern University, Shenyang, China.*
- *R. Duits is with the department of Mathematics and Computer Science, TU/e; he is also affiliated to the department of Biomedical Enginering, TU/e. Email: r.duits@tue.nl*

Manuscript received ...., ...; revised ...., ...

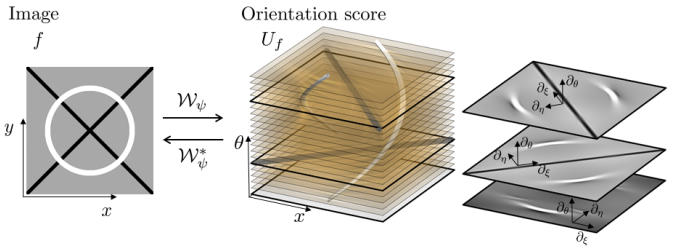

Fig. 2. In orientation scores U_f, constructed from an image f via the orientation score transform $\mathcal{W}_\psi$, we make use of a left-invariant derivative frame $\{\partial_\xi, \partial_\eta, \partial_\theta\}$ that is aligned with the orientation θ corresponding to each layer in the score. Three slices and the corresponding left-invariant frames are shown separately (at $\theta \in \{0, \frac{\pi}{4}, \frac{3\pi}{4}\}$).

name *ridge regression* [8], when taking

$$S(t) = \sum_{i=1}^{N} \left((t, f_i)_{\mathbb{L}_2(\mathbb{R}^2)} - y_i \right)^2, \quad \text{and} \quad R(t) = \mu \|t\|^2_{\mathbb{L}_2(\mathbb{R}^2)},$$

where f_i is one of N image patches, $y_i \in \{0, 1\}$ is the corresponding desired filter response, and where μ is a parameter weighting the regularization term. The regression is then from an input image patch f_i to a desired response y_i, and the template t can be regarded as the "set of weights" that are optimized in the regression problem. In this article we consider both quadratic (linear regression) and logistic (logistic regression) losses in S. For regularization we consider terms of the form

$$R(t) = \lambda \int_{\mathbb{R}^2} \|\nabla t(\mathbf{x})\|^2 \mathrm{d}\mathbf{x} + \mu \|t\|^2_{\mathbb{L}_2(\mathbb{R}^2)},$$

and thus combine the classical ridge regression with a smoothing term (weighted by λ).

In our extension of smoothed regression to orientation scores we employ similar techniques. However, here we must recognize a curved geometry on the domain $\mathbb{R}^2 \rtimes S^1$, which we identify with the group of roto-translations: the Lie group $SE(2)$ equipped with group product

$$g \cdot g' = (\mathbf{x}, \theta) \cdot (\mathbf{x}', \theta') = (\mathbf{R}_\theta \mathbf{x}' + \mathbf{x}, \theta + \theta'). \tag{2}$$

In this product the orientation θ influences the product on the spatial part. Therefore we write $\mathbb{R}^2 \rtimes S^1$ instead of $\mathbb{R}^2 \times S^1$, as it is a semi-direct group product (and not a direct product). Accordingly, we must work with a rotating derivative frame (instead of axis aligned derivatives) that is aligned with the group elements $(\mathbf{x}, \theta) \in SE(2)$, see e.g. the $(\partial_\xi, \partial_\eta, \partial_\theta)$-frames in Fig. 2. This derivative frame allows for (anisotropic) smoothing along oriented structures. As we will show in this article (Sec. A), the proposed smoothing scheme has the probabilistic interpretation of time integrated Brownian motion on $SE(2)$ [9], [10].

Regression and Group Theory. Regularization in (generalized) linear regression generally leads to more robust classifiers/regressions, especially when a low number of training samples are available. Different types of regularizations in regression problems have been intensively studied in e.g. [11], [12], [13], [14], [15], and the choice for regularization-type depends on the problem: E.g. $\mathbb{L}_1$-type regularization is often used to sparsify the regression weights, whereas $\mathbb{L}_2$-type regularization is more generally used to prevent overfitting by penalizing outliers (e.g. in ridge regression [8]). Smoothing of regression coefficients by penalizing the $\mathbb{L}_2$-norm of the derivative along the coefficients is less common, but it can have a significant effect on performance [13], [16].

We solve problem (1) in the context of smoothing splines: We discretize the problem by expanding the templates in a finite B-spline basis, and optimize over the spline coefficients. For d-dimensional Euclidean spaces, smoothing splines have been well studied [17], [18], [19], [20]. In this paper, we extend the concept to the curved space $SE(2)$ and provide explicit forms of the discrete regularization matrices. Furthermore, we show that the extended framework can be used for time integrated Brownian motions on $SE(2)$, and show near to perfect comparisons to the exact solutions found in [9], [10].

In general, statistics and regression on Riemannian manifolds are powerful tools in medical imaging and computer vision [21], [22], [23], [24]. More specifically in pattern matching and registration problems Lie groups are often used to describe deformations. E.g. in [25] the authors learn a regression function $\mathbb{R}^m \to A(2)$ from a discrete m-dimensional feature vector to a deformation in the affine group $A(2)$. Their purpose is object tracking in video sequences. This work is however not concerned with deformation analysis, we instead learn a regression function $\mathbb{L}_2(SE(2)) \to \mathbb{R}$ from continuous densities on the Lie group $SE(2)$ (obtained via an invertible orientation score transform) to a desired filter response. Our purpose is object detection in 2D images. In our regression we impose smoothed regression with a time-integrated hypo-elliptic Brownian motion prior and thereby extend least squares regression to smoothed regression on SE(2) involving first order variation in Sobolev-type of norms.

Application Area of the Proposed Method. The strength of our approach is demonstrated with the application to anatomical landmark detection in medical retinal images and pupil localization in regular camera images. In the retinal application we consider the problem of detecting the optic nerve head (ONH) and the fovea. Many image analysis applications require the robust, accurate and fast detection of these structures, see e.g. [26], [27], [28], [29]. In all three detection problems the objects of interest are characterized by (surrounding) curvilinear structures (blood vessels in the retina; eyebrow, eyelid, pupil and other contours for pupil detection), which are conveniently represented in *invertible* orientation scores. The invertibility condition implies that all image data is contained in the orientation score [30] [7]. With the proposed method we achieve state-of-the-art results both in terms of detection performance and speed: high detection performance is achieved by learning templates that make optimal use of the line patterns in orientation scores; speed is achieved by a simple, yet effective, cross-correlation template matching approach.

Contribution of this Work. This article builds upon two published conference papers [31], [32]. In the first we demonstrated that high detection performance could be achieved by considering cross-correlation based template matching in $SE(2)$, using only handcrafted templates and with the application of ONH detection in retinal images [31]. In the second we then showed on the same application that better performance could be achieved by training templates using the optimization of energy functionals of the form of

(1), where then only a (left-invariant) smoothing regularizer was considered [32]. In this article we provide a complete framework for training of templates and matching on $SE(2)$ and contribute to literature by:

1) Extending the linear regression $SE(2)$ framework [32] to logistic regression, with clear benefits in pupil detection using a single template (with an increase of success rate from 76% to 94%).
2) Studying different types of regression priors, now introducing also a ridge regression prior.
3) We show that the $SE(2)$ smoothing prior corresponds to time-integrated hypo-elliptic diffusion on $SE(2)$, providing a Brownian motion interpretation.
4) We show the generic applicability of our method: with the exact same settings of our algorithm we obtain state-of-the-art results on three different applications (ONH detection, cf. Ch. 4.2 and Table 2, fovea detection, cf. Subsec. 4.3 and Table 3, and pupil detection, cf. Subsec. 4.4 and Fig. 5).
5) Improving previous results on ONH detection (reducing the number of failed detections to 3 out of 1737 images).
6) Making our code publicly available at http://erikbekkers.bitbucket.org/TMSE2.html.

Paper Outline. The remainder of this paper is organized as follows. In Sec. 2 we provide the theory for template matching and template construction in the $\mathbb{R}^2$-case. The theory is then extended to the $SE(2)$-case in Sec. 3. Additionally, in Sec. A we provide a probabilistic interpretation of the proposed $SE(2)$ prior, and relate it to Brownian motions on $SE(2)$. In Sec. 4 we apply the method to retinal images for ONH (Subsec. 4.2) and fovea detection (Subsec. 4.3), and to regular camera images for pupil detection (Subsec. 4.4). Finally, we conclude the paper in Sec. 5.

2 Template Matching & Regression on $\mathbb{R}^2$

2.1 Object Detection via Cross-Correlation

We are considering the problem of finding the location of objects (with specific orientation patterns) in an image. While in principle an image may contain multiple objects of interest, the applications discussed in this paper only require the detection of one object per image. We search for the most likely location

$$\mathbf{x}^* = \underset{\mathbf{x}\in\mathbb{R}^2}{\operatorname{argmax}} \; P(\mathbf{x}), \tag{3}$$

with $P(\mathbf{x}) \in \mathbb{R}$ denoting the objective functional for finding the object of interest at location $\mathbf{x}$. We define P based on inner products in a *linear regression* and *logistic regression* context, where we respectively define P by

$$P(\mathbf{x}) = P^{\mathbb{R}^2}_{lin}(\mathbf{x}) := (\mathcal{T}_{\mathbf{x}}\, t, f)_{\mathbb{L}_2(\mathbb{R}^2)}, \tag{4}$$

or

$$\begin{aligned} &P(\mathbf{x}) = P^{\mathbb{R}^2}_{log}(\mathbf{x}) := \sigma\left((\mathcal{T}_{\mathbf{x}}\, t, f)_{\mathbb{L}_2(\mathbb{R}^2)}\right), \\ &\text{with} \quad \sigma(x) = e^x/(1+e^x), \end{aligned} \tag{5}$$

where $\mathcal{T}_{\mathbf{x}}$ denotes translation by $\mathbf{x}$ via

$$(\mathcal{T}_{\mathbf{x}} t)(\tilde{\mathbf{x}}) = t(\tilde{\mathbf{x}} - \mathbf{x}),$$

and where the $\mathbb{L}_2(\mathbb{R}^2)$ inner product is given by

$$(t, f)_{\mathbb{L}_2(\mathbb{R}^2)} := \int_{\mathbb{R}^2} \overline{t(\tilde{\mathbf{x}})} f(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}, \tag{6}$$

with associated norm $\|\cdot\|_{\mathbb{L}_2(\mathbb{R}^2)} = \sqrt{(\cdot,\cdot)_{\mathbb{L}_2(\mathbb{R}^2)}}$. Note that the inner-product based potentials $P(\mathbf{x})$ can be efficiently evaluated for each $\mathbf{x}$ using convolutions.

For a generalization of cross-correlation based template matching to *normalized* cross correlation, we refer the reader to the supplementary materials. For speed considerations we will however not use normalized cross correlation, but instead use a (fast) preprocessing step to locally normalize the images (cf. Subsec. 4.2.1).

2.2 Optimizing t Using Linear Regression

Our aim is to construct templates t that are "aligned" with image patches that contain the object of interest, and which are orthogonal to non-object patches. Hence, template t is found via the minimization of the following energy

$$\begin{aligned} E_{lin}(t) = \sum_{i=1}^{N} \left((t, f_i)_{\mathbb{L}_2(\mathbb{R}^2)} - y_i\right)^2 \\ + \lambda \int_{\mathbb{R}^2} \|\nabla t(\tilde{\mathbf{x}})\|^2 \mathrm{d}\tilde{\mathbf{x}} + \mu \, \|t\|^2_{\mathbb{L}_2(\mathbb{R}^2)}, \end{aligned} \tag{7}$$

with f_i one of the N training patches extracted from an image $f_{\mathbf{x}}$, and y_i the corresponding label ($y_i = 1$ for *objects* and $y_i = 0$ for *non-objects*). In (7), the data-term (first term) aims for alignment of template t with object patches, in which case the inner product $(t, f_i)_{\mathbb{L}_2(\mathbb{R}^2)}$ is ideally one, and indeed aims orthogonality to non-object patches (in which case the inner product is zero). The second term enforces spatial smoothness of the template by penalizing its gradient, controlled by λ. The third (ridge) term improves stability by dampening the $\mathbb{L}_2$-norm of t, controlled by μ.

2.3 Optimizing t Using Logistic Regression

In object detection we are essentially considering a two-class classification problem: the object is either present or it is not. In this respect, the quadratic loss term in (7) might not be the best choice as it penalizes any deviation from the desired response y_i, regardless of whether or not the response $(t, f_i)_{\mathbb{L}_2(\mathbb{R}^2)}$ is on the correct side of a decision boundary. In other words, the aim is not necessarily to construct a template that best maps an image patch f_i to a response $y_i \in \{0, 1\}$, but rather the aim is to construct a template that best makes the separation between *object* and *non-object* patches. With this in mind we resort to the logistic regression model, in which case we interpret the non-linear objective functional given in (5) as a probability, and define

$$\begin{aligned} &p_1(f_i \,;\, t) \;\; = p(f_i \,;\, t), \\ &p_0(f_i \,;\, t) \;\; = 1 - p(f_i \,;\, t), \\ &\text{with } p(f_i \,;\, t) = \sigma\left((t, f_i)_{\mathbb{L}_2(\mathbb{R}^2)}\right), \end{aligned} \tag{8}$$

with $p_1(f_i; t)$ and $p_0(f_i; t)$ denoting respectively the probabilities of a patch f_i being an *object* or *non-object* patch. Our aim is now to maximize the likelihood (of each patch f_i having maximum probability $p_{y_i}(f_i; t)$ for correct label y_i):

$$\ell(t) = \prod_{i=1}^{N} p_{y_i}(f_i; t) = \prod_{i=1}^{N} p(f_i; t)^{y_i} (1 - p(f_i; t))^{1-y_i}. \tag{9}$$

We maximize the log-likelood instead, which is given by

$$\begin{aligned}\ell_{log}(t) &:= \log(\ \ell(t)\ ) \\ &= \sum_{i=1}^{N} \log(\ p(f_i;t)^{y_i}(1-p(f_i;t))^{1-y_i}\ ) \\ &= \sum_{i=1}^{N} y_i (t,f_i)_{\mathbb{L}_2(\mathbb{R}^2)} - \log\left(1+e^{(t,f_i)_{\mathbb{L}_2(\mathbb{R}^2)}}\right).\end{aligned} \tag{10}$$

Maximizing (10) is known as the problem of logistic regression. Similar to the linear regression case, we impose additional regularization and define the following regularized logistic regression energy, which we aim to *maximize*:

$$E^{\ell}_{log}(t) = \ell_{log}(t) - \lambda \int_{\mathbb{R}^2} \|\nabla t(\tilde{\mathbf{x}})\|^2 \mathrm{d}\tilde{\mathbf{x}} - \ \mu\ \|t\|^2_{\mathbb{L}_2(\mathbb{R}^2)}. \tag{11}$$

2.4 Template Optimization in a B-Spline Basis

Templates in a B-Spline Basis. In order to solve the optimizations (7) and (11), the template is described in a basis of direct products of n-th order B-splines B^n:

$$t(x,y) = \sum_{k=1}^{N_k}\sum_{l=1}^{N_l} c_{k,l}\ B^n\left(\frac{x}{s_k}-k\right) B^n\left(\frac{y}{s_l}-l\right), \tag{12}$$

with $B^n(x) = \left(1_{\left[-\frac{1}{2},\frac{1}{2}\right]} *^{(n)} 1_{\left[-\frac{1}{2},\frac{1}{2}\right]}\right)(x)$ a n-th order B-spline obtained by n-fold convolution of the indicator function $1_{\left[-\frac{1}{2},\frac{1}{2}\right]}$, and $c_{k,l}$ the coefficients belonging to the shifted B-splines. Here s_k and s_l scale the B-splines and typically depend on the number N_k and N_l of B-splines.

Linear Regression. By substitution of (12) in (7), the energy functional can be expressed in matrix-vector form (see Section 2 of the supplementary materials):

$$E^{B}_{lin}(\mathbf{c}) = \|S\mathbf{c}-\mathbf{y}\|^2 + \lambda\ \mathbf{c}^{\dagger} R\mathbf{c} + \mu\ \mathbf{c}^{\dagger} I\mathbf{c}. \tag{13}$$

Regarding our notations we note that for spatial template t given by (12) we have $E_{lin}(t) = E^{B}_{lin}(\mathbf{c})$, and label 'B' indicates finite expansion in the B-spline basis. The minimizer of (13) is given by

$$(S^{\dagger}S + \lambda R + \mu I)\mathbf{c} = S^{\dagger}\mathbf{y}, \tag{14}$$

with † denoting the conjugate transpose, and I denoting the identity matrix. Here S is a $[N \times N_k N_l]$ matrix given by

$$\begin{aligned} S &= \{(s^i_{1,1},...,s^i_{1,N_l},s^i_{2,1},...,s^i_{2,N_l},...,...,s^i_{N_k,N_l})\}_{i=1}^{N}, \\ s_{k,l} &= (\ B^n_{s_k s_l} * f_i\ )(k,l), \end{aligned} \tag{15}$$

with $B^n_{s_k s_l}(x,y) = B^n\left(\frac{x}{s_k}\right) B^n\left(\frac{y}{s_l}\right)$, for all (x,y) on the discrete spatial grid on which the input image $f_D : \{1,N_x\} \times \{1,N_y\} \to \mathbb{R}$ is defined. Here N_k and N_l denote the number of splines in resp. x and y direction, and $s_k = \frac{N_x}{N_k}$ and $s_l = \frac{N_y}{N_l}$ are the corresponding resolution parameters. The $[N_k N_l \times 1]$ column vector $\mathbf{c}$ contains the B-spline coefficients, and the $[N \times 1]$ column vector $\mathbf{y}$ contains the labels, stored in the following form

$$\begin{aligned} \mathbf{c} &= (c_{1,1},...,c_{1,N_l},c_{2,1},...,c_{2,N_l},...,...,c_{N_k,N_l})^T \\ \mathbf{y} &= (y_1,y_2,...,y_N)^T. \end{aligned} \tag{16}$$

The $[N_k N_l \times N_k N_l]$ regularization matrix R is given by

$$R = R^{s_k}_x \otimes R^{s_l}_x + R^{s_k}_y \otimes R^{s_l}_y, \tag{17}$$

where $\otimes$ denotes the Kronecker product, and with

$$\begin{aligned} R^{s_k}_x(k,k') &= -\tfrac{1}{s_k}\tfrac{\partial^2 B^{2n+1}}{\partial x^2}(k'-k), \\ R^{s_l}_x(l,l') &= s_l B^{2n+1}(l'-l), \\ R^{s_k}_y(k,k') &= s_k B^{2n+1}(k'-k), \\ R^{s_l}_y(l,l') &= -\tfrac{1}{s_l}\tfrac{\partial^2 B^{2n+1}}{\partial y^2}(l'-l), \end{aligned} \tag{18}$$

with $k,k' = 1,2,...,N_k$ and $l,l' = 1,2,...,N_l$. The coefficients $\mathbf{c}$ can then be computed by solving (14) directly, or via linear system solvers such as conjugate gradient descent. For a derivation of the regularization matrix R we refer to supplementary materials, Sec. 2.

Logistic Regression. The logistic regression log-likelihood functional (11) can be expressed in matrix-vector notations as follows:

$$\begin{aligned} E^{\ell,B}_{log}(\mathbf{c}) = &\left[\mathbf{y}^{\dagger} S\mathbf{c} - \mathbf{1}^{\dagger}_N \log(\mathbf{1}_N + \exp(S\mathbf{c}))\right] \\ &- \lambda\ \mathbf{c}^{\dagger} R\mathbf{c} - \mu\ \mathbf{c}^{\dagger} I\mathbf{c}, \end{aligned} \tag{19}$$

where $\mathbf{1}_N = \{1,1,...,1\}^T \in \mathbb{R}^{N\times 1}$, and where the exponential and logarithm are evaluated element-wise. We follow a standard approach for the optimization of (19), see e.g. [11], and find the minimizer by settings the derivative to $\mathbf{c}$ to zero

$$\nabla_{\mathbf{c}} E^{\ell,B}_{log}(\mathbf{c}) = S^T(\mathbf{y}-\mathbf{p}) - \lambda\ R\mathbf{c} - \mu\ I\mathbf{c} = \mathbf{0}, \tag{20}$$

with $\mathbf{p} = (p_1,...,p_N)^T \in \mathbb{R}^{N\times 1}$, with $p_i = \sigma((S\mathbf{c})_i)$. To solve (20), we use a Newton-Raphson optimization scheme. This requires computation of the Hessian matrix, given by

$$\mathcal{H}(E^{\ell,B}_{log}) = -(S^T W S + \lambda\ R + \mu\ I), \tag{21}$$

with diagonal matrix $W = \underset{i\in\{1,...,N\}}{\mathrm{diag}}\{p_i(1-p_i)\}$. The Newton-Raphson update rule is then given by

$$\begin{aligned} \mathbf{c}^{new} &= \mathbf{c}^{old} - \mathcal{H}(E^{\ell,D}_{log})^{-1}(\nabla_{\mathbf{c}} E^{\ell,D}_{log}(\mathbf{c})) \\ &= (S^T W S + \lambda\ R + \mu\ I)^{-1} S^T W\mathbf{z}, \end{aligned} \tag{22}$$

with $\mathbf{z} = S\mathbf{c}^{old} + W^{-1}(\mathbf{y}-\mathbf{p})$, see e.g. [11, ch. 4.4]. Optimal coefficients found at convergence are denoted with $\mathbf{c}^*$.

Summarizing, we obtain the solution of (3) by substituting the optimized B-spline coefficients $\mathbf{c}^*$ into (12), and the resulting t enters (4) or (5). The most likely object location $\mathbf{x}^*$ is then found via (3).

3 Template Matching & Regression on $SE(2)$

This section starts with details on the representation of image data in the form of orientation scores (Subsec. 3.1). Then, we repeat the sections from Sec. 2 in Subsections 3.2 to 3.5, but now in the context of the extended domain $SE(2)$.

3.1 Orientation Scores on $SE(2)$

Transformation. An orientation score, constructed from image $f : \mathbb{R}^2 \to \mathbb{R}$, is defined as a function $U_f : \mathbb{R}^2 \rtimes S^1 \to \mathbb{C}$ and depends on two variables $(\mathbf{x},\theta)$, where $\mathbf{x} = (x,y) \in \mathbb{R}^2$ denotes position and $\theta \in [0,2\pi)$ denotes the orientation variable. An orientation score U_f of image f can be constructed by means of correlation with some anisotropic wavelet ψ via

$$U_f(\mathbf{x},\theta) = (\mathcal{W}_{\psi} f)(\mathbf{x},\theta) = \int_{\mathbb{R}^2} \overline{\psi(\mathbf{R}^{-1}_{\theta}(\tilde{\mathbf{x}}-\mathbf{x}))} f(\tilde{\mathbf{x}})\mathrm{d}\tilde{\mathbf{x}}, \tag{23}$$

where $\psi \in \mathbb{L}_2(\mathbb{R}^2)$ is the correlation kernel, aligned with the x-axis, where $\mathcal{W}_\psi$ denotes the transformation between image f and orientation score U_f, $\psi_\theta(\mathbf{x}) = \psi(\mathbf{R}_\theta^{-1}\mathbf{x})$, and $\mathbf{R}_\theta$ is a counter clockwise rotation over angle θ.

In this work we choose cake wavelets [6], [7] for ψ. While in general any kind of anisotropic wavelet could be used to lift the image to $SE(2)$, cake wavelets ensure that no data-evidence is lost during the transformation: By design the set of all rotated wavelets uniformly cover the full Fourier domain of disk-limited functions with zero mean, and have thereby the advantage over other oriented wavelets (s.a. Gabor wavelets for specific scales) that they capture all scales and allow for a stable inverse transformation $\mathcal{W}_\psi^*$ from the score back to the image [6], [10].

Left-Invariant Derivatives. The domain of an orientation score is essentially the classical Euclidean motion group $SE(2)$ of planar translations and rotations, and is equipped with group product $g \cdot g' = (\mathbf{x}, \theta) \cdot (\mathbf{x}', \theta') = (\mathbf{R}_\theta \mathbf{x}' + \mathbf{x}, \theta + \theta')$. Here, we can recognize a curved geometry (cf. Fig. 2), and it is therefore useful to work in rotating frame of reference. As such, we use the left invariant derivative frame [9], [10]:

$$\{\partial_\xi := \cos\theta\, \partial_x + \sin\theta\, \partial_y, \partial_\eta := -\sin\theta\, \partial_x + \cos\theta\, \partial_y, \partial_\theta\}. \tag{24}$$

Using this derivative frame we will construct in Subsec. 3.3 a regularization term in which we can control the amount of (anisotropic) smoothness along line structures.

3.2 Object Detection via Cross-Correlation

As in Section 2, we search for the most likely object location $\mathbf{x}^*$ via (3), but now we define functional P respectively for the linear and logistic regression case in $SE(2)$ by:

$$P(\mathbf{x}) = P_{lin}^{SE(2)}(\mathbf{x}) := (\mathcal{T}_\mathbf{x}\, T, U_f)_{\mathbb{L}_2(SE(2))}, \quad \text{or} \tag{25}$$

$$P(\mathbf{x}) = P_{log}^{SE(2)}(\mathbf{x}) := \sigma\left((\mathcal{T}_\mathbf{x}\, T, U_f)_{\mathbb{L}_2(SE(2))}\right), \tag{26}$$

with $(\mathcal{T}_\mathbf{x} T)(\tilde{\mathbf{x}}, \tilde{\theta}) = T(\tilde{\mathbf{x}} - \mathbf{x}, \tilde{\theta})$. The $\mathbb{L}_2(SE(2))$-inner product is defined by

$$(T, U_f)_{\mathbb{L}_2(SE(2))} := \int_{\mathbb{R}^2} \int_0^{2\pi} \overline{T(\tilde{\mathbf{x}}, \tilde{\theta})} U_f(\tilde{\mathbf{x}}, \tilde{\theta}) \mathrm{d}\tilde{\mathbf{x}} \mathrm{d}\tilde{\theta}, \tag{27}$$

with norm $\|\cdot\|_{\mathbb{L}_2(SE(2))} = \sqrt{(\cdot,\cdot)_{\mathbb{L}_2(SE(2))}}$.

3.3 Optimizing T Using Linear Regression

Following the same reasoning as in Section 2.2 we search for the template that minimizes

$$\mathcal{E}_{lin}(T) = \sum_{i=1}^{N} \left((T, U_{f_i})_{\mathbb{L}_2(SE(2))} - y_i\right)^2 + \lambda \int_{\mathbb{R}^2} \int_0^{2\pi} \|\nabla T(\tilde{\mathbf{x}}, \tilde{\theta})\|_D^2 \mathrm{d}\tilde{\mathbf{x}} \mathrm{d}\tilde{\theta} + \mu \|T\|^2_{\mathbb{L}_2(SE(2))}, \tag{28}$$

with smoothing term:

$$\|\nabla T(g)\|_D^2 = D_{\xi\xi} \left|\frac{\partial T}{\partial \xi}(g)\right|^2 + D_{\eta\eta} \left|\frac{\partial T}{\partial \eta}(g)\right|^2 + D_{\theta\theta} \left|\frac{\partial T}{\partial \theta}(g)\right|^2. \tag{29}$$

Here, $\nabla T = (\frac{\partial T}{\partial \xi}, \frac{\partial T}{\partial \eta}, \frac{\partial T}{\partial \theta})^T$ denotes the left-invariant gradient. Note that ∂_ξ gives the spatial derivative in the direction aligned with the orientation score kernel used at layer θ, recall Fig. 2. The parameters $D_{\xi\xi}$, $D_{\eta\eta}$ and $D_{\theta\theta} \geq 0$ are then used to balance the regularization in the three directions. Similar to this problem, first order Tikhonov-regularization on $SE(2)$ is related, via temporal Laplace transforms, to left–invariant diffusions on the group $SE(2)$ (Sec. A), in which case $D_{\xi\xi}$, $D_{\eta\eta}$ and $D_{\theta\theta}$ denote the diffusion constants in ξ, η and θ direction. Here we set $D_{\xi\xi} = 1$, $D_{\eta\eta} = 0$, and thereby we get Laplace transforms of hypo-elliptic diffusion processes [10], [33]. Parameter $D_{\theta\theta}$ can be used to tune between isotropic (large $D_{\theta\theta}$) and anisotropic (low $D_{\theta\theta}$) diffusion (see e.g. [32, Fig. 3]). Note that anisotropic diffusion, via a low $D_{\theta\theta}$, is preferred as we want to maintain line structures in orientation scores.

3.4 Optimizing T Using Logistic Regression

Similarly to what is done in Subsec. 2.3 we can change the quadratic loss of (28) to a logistic loss, yielding the following energy functional

$$\mathcal{E}_{log}(T) = \mathscr{L}_{log}(T) - \lambda \int_{\mathbb{R}^2} \int_0^{2\pi} \|\nabla T(\tilde{\mathbf{x}}, \tilde{\theta})\|_D^2 \mathrm{d}\tilde{\mathbf{x}} \mathrm{d}\tilde{\theta} - \mu \|T\|^2_{\mathbb{L}_2(SE(2))}, \tag{30}$$

with log-likelihood (akin to (10) for the $\mathbb{R}^2$ case)

$$\mathscr{L}_{log}(T) = \sum_{i=1}^{N} y_i (T, U_{f_i})_{\mathbb{L}_2(SE(2))} - \log\left(1 + e^{(T, U_{f_i})_{\mathbb{L}_2(SE(2))}}\right). \tag{31}$$

The optimization of (28) and (30) follows quite closely the procedure as described in Sec. 2 for the 2D case. In fact, when T is expanded in a B-spline basis, the exact same matrix-vector formulation can be used.

3.5 Template Optimization in a B-Spline Basis

Templates in a B-Spline Basis. The template T is expanded in a B-spline basis as follows

$$T(x, y, \theta) = \sum_{k=1}^{N_k} \sum_{l=1}^{N_l} \sum_{m=1}^{N_m} c_{k,l,m} \cdot B^n\left(\frac{x}{s_k} - k\right) B^n\left(\frac{y}{s_l} - l\right) B^n\left(\frac{\theta \bmod 2\pi}{s_m} - m\right), \tag{32}$$

with N_k, N_l and N_m the number of B-splines in respectively the x, y and θ direction, $c_{k,l,m}$ the corresponding basis coefficients, and with angular resolution parameter $s_m = 2\pi / N_m$.

Linear Regression. The shape of the minimizer of energy functional $\mathcal{E}_{lin}(T)$ in the $SE(2)$ case is the same as for $E_{lin}(t)$ in the $\mathbb{R}^2$ case, and is again of the form given in (13). However, now the definitions of S, R and $\mathbf{c}$ are different. Now, S is a $[N \times N_k N_l N_m]$ matrix given by

$$S = \{(s^i_{1,1,1}, ..., s^i_{1,1,N_m}, ..., s_{1,N_l,N_m}, ..., s^i_{N_k,N_l,N_m})\}_{i=1}^N, \quad s_{k,l,m} = (\, B^n_{s_k s_l s_m} * U_{f_i}\,)(k,l,m), \tag{33}$$

with $B^n_{s_k s_l s_m}(x, y, \theta) = B^n\left(\frac{x}{s_k}\right) B^n\left(\frac{y}{s_l}\right) B^n\left(\frac{\theta \bmod 2\pi}{s_m}\right)$. Vector $\mathbf{c}$ is a $[N_k N_l N_m \times 1]$ column vector containing the B-spline coefficients and is stored as follows:

$$\mathbf{c} = (c_{1,1,1}, ..., c_{1,1,N_m}, ..., c_{1,N_l,N_m}, ..., c_{N_k,N_l,N_m})^T. \tag{34}$$

The explicit expression and the derivation of $[N_k N_l N_m \times N_k N_l N_m]$ matrix R, which encodes the left invariant derivatives, can be found in the supplementary materials Sec. 2.

Logistic Regression Also for the logistic regression case we optimize energy functional (30) in the same form as (11) in the $\mathbb{R}^2$ case, by using the corresponding expressions for S, R, and $\mathbf{c}$ in Eq. (19). These expressions can be inserted in the functional (19) and again the same techniques (as presented in Subsection 2.4) can be used to minimize this cost on $SE(2)$.

3.6 Probabilistic Interpretation of the $SE(2)$ Smoothing Prior

In this section we only provide a brief introduction to the probabilistic interpretation of the $SE(2)$ smoothing prior, and refer the interested reader to the supplementary materials for full details. Consider the classic approach to noise suppression in images via diffusion regularizations with PDE's of the form

$$\begin{cases} \frac{\partial}{\partial \tau} u &= \Delta u, \\ u|_{\tau=0} &= u_0, \end{cases} \tag{35}$$

where Δ denotes the Laplace operator. Solving (40) for any diffusion time $\tau > 0$ gives a smoothed version of the input u_0. The time-resolvent process of the PDE is defined by the Laplace transform with respect to τ; time τ is integrated out using a memoryless negative exponential distribution $P(\mathcal{T} = \tau) = \alpha e^{-\alpha \tau}$. Then, the time integrated solutions

$$t(\mathbf{x}) = \alpha \int_0^\infty u(\mathbf{x}, \tau) e^{-\alpha \tau} \mathrm{d}\tau,$$

with decay parameter α, are in fact the solutions [34]

$$t = \underset{t \in \mathbb{L}_2(\mathbb{R}^2)}{\operatorname{argmin}} \left[\|t - t_0\|^2_{\mathbb{L}_2(\mathbb{R}^2)} + \lambda \int_{\mathbb{R}^2} \|\nabla t(\tilde{\mathbf{x}})\|^2 \, \mathrm{d}\tilde{\mathbf{x}} \right], \tag{36}$$

with $\lambda = \alpha^{-1}$. Such time integrated diffusions (Eq. (36)) can also be obtained by optimization of the linear regression functionals given by Eq. (7) and Eq. (25) for the $\mathbb{R}^2$ and $SE(2)$ case respectively.

In the supplementary materials we establish this connection for the $SE(2)$ case, and show how the smoothing regularizer in (28) and (30) relates to Laplace transforms of hypo-elliptic diffusions on the group $SE(2)$ [9], [10]. More precisely, we formulate a special case of our problem (the *single patch problem*) which involves only a single training sample U_{f_1}, and show in a formal theorem that the solution is up to scalar multiplication the same as the resolvent hypo-elliptic diffusion kernel. The underlying probabilistic interpretation is that of Brownian motions on $SE(2)$, where the resolvent hypo-elliptic diffusion kernel gives a probability density of finding a random brush stroke at location $\mathbf{x}$ with orientation θ, given that a 'drunkman's pencil' starts at the origin at time zero.

In the supplementary materials we demonstrate the high accuracy of our discrete numeric regression method using B-spline expansions with near to perfect comparisons to the continuous exact solutions of the single patch problem. In fact, we have established an efficient B-spline finite element implementation of hypo-elliptic Brownian motions on $SE(2)$, in addition to other numerical approaches in [9].

4 Applications

Our applications of interest are in retinal image analysis. In this section we establish and validate an algorithm pipeline for the detection of the optic nerve head (Subsec. 4.2) and fovea (Subsec. 4.3) in retinal images, and the pupil (Subsec. 4.4) in regular camera images. Before we proceed to the application sections, we first describe the experimental set-up (Subsec. 4.1). All experiments discussed in this section are reproducible; the data (with annotations) as well as the full code (Wolfram *Mathematica* notebooks) used in the experiments are made available at: http://erikbekkers.bitbucket.org/TMSE2.html. In the upcoming sections we only report the most relevant experimental results. More details on each application (examples of training samples, implementation details, a discussion on parameter settings, computation times, and examples of successful/failed detections) are provided in the supplementary materials.

4.1 The experimental set-up

Templates. In our experiments we compare the performance of different template types, which we label as follows:

- A: Templates obtained by taking the average of all positive patches ($y_i = 1$) in the training set, then normalized to zero mean and unit standard deviation.
- B: Templates optimized without any regularization.
- C: Templates optimized with an optimal μ, and with $\lambda = 0$.
- D: Templates optimized with an optimal λ and with $\mu = 0$.
- E: Templates optimized with optimal μ and λ.

The trained templates (B-E) are obtained either via linear or logistic regression in the $\mathbb{R}^2$ setting (see Subsec. 2.4 and Subsec. 2.4), or in the $SE(2)$ setting (see Subsec. 3.5 and Subsec. 3.5). In both the $\mathbb{R}^2$ and $SE(2)$ case, linear regression based templates are indicated with subscript $_{lin}$, and logistic regression based templates with $_{log}$. Optimality of parameter values is defined using generalized cross validation (GCV), which we soon explain in this section. We generally found that (via optimization using GCV) the optimal settings for template E were $\mu \approx 0.5\mu^*$, and $\lambda \approx 0.5\lambda^*$, with μ^* and λ^* respectively the optimal parameters for template C and D.

Matching with Multiple Templates. When performing template matching, we use Eq. (4) and Eq. (25) for respectively the $\mathbb{R}^2$ and $SE(2)$ case for templates obtained via linear regression and for template A. For templates obtained via logistic regression we use respectively Eq. (5) and Eq. (26). When we combine multiple templates we simply add the objective functionals. E.g, when combining template $C_{lin:\mathbb{R}^2}$ and $D_{log:SE(2)}$ we solve the problem

$$\mathbf{x}^* = \underset{\mathbf{x} \in \mathbb{R}^2}{\operatorname{argmax}} \; P^{\mathbb{R}^2}_{C_{lin}}(\mathbf{x}) + P^{SE(2)}_{D_{log}}(\mathbf{x}),$$

where $P^{\mathbb{R}^2}_{C_{lin}}(\mathbf{x})$ is the objective functional (see Eq. (4)) obtained with template $C_{lin:\mathbb{R}^2}$, and $P^{SE(2)}_{D_{log}}(\mathbf{x})$ (see Eq. (26)) is obtained with template $D_{log:SE(2)}$.

Rotation and Scale Invariance. The proposed template matching scheme can adapted for rotation-scale invariant matching, this is discussed in Sec. 5 of the supplementary materials. For a generic object recognition task, however,

global rotation or scale invariance are not necessarily desired properties. Datasets often contain objects in a human environment context, in which some objects tend to appear in specific orientations (e.g. eye-brows are often horizontal above the eye and vascular trees in the retina depart the ONH typically along a vertical axis). Discarding such knowledge by introducing rotation/scale invariance is likely to have an adversary effect on the performance, while increasing computational load. In Sec. 5 of the supplementary materials we tested a rotation/scale invariant adaptation of our method and show that in the three discussed applications this did indeed not lead to improved results, but in fact worsened the results slightly.

Automatic Parameter Selection via Generalized Cross Validation. An ideal template generalizes well to new data samples, meaning that it has low prediction error on independent data samples. One method to predict how well the system generalizes to new data is via generalized cross validation (GCV), which is essentially an approximation of leave-one-out cross validation [35]. The vector containing all predictions is given by $\tilde{\mathbf{y}} = S\mathbf{c}_{\mu,\lambda}$, in which we can substitute the solution for $\mathbf{c}_{\mu,\lambda}$ (from Eq. (14)) to obtain

$$\begin{aligned} \tilde{\mathbf{y}} &= A_{\mu,\lambda}\mathbf{y}, \qquad \text{with} \\ A_{\mu,\lambda} &= S(S^\dagger S + \lambda R + \mu I)^{-1} S^\dagger, \end{aligned} \tag{37}$$

where $A_{\mu,\lambda}$ is the so-called *smoother matrix*. Then the generalized cross validation value [35] is defined as

$$GCV(\mu,\lambda) \equiv \frac{\frac{1}{N}\|\Omega(I - A_{\mu,\lambda})\mathbf{y}\|^2}{\left(1 - \mathrm{trace}(A_{\mu,\lambda})/N\right)^2}. \tag{38}$$

In the retinal imaging applications we set $\Omega = I$. In the pupil detection application we set $\Omega = \underset{i\in\{1,\ldots,N\}}{\mathrm{diag}}\{y_i\}$. As such, we do not penalize errors on negative samples as here the diversity of negative patches is too large for parameter optimization via GCV. Parameter settings are considered optimal when they minimize the GCV value.

In literature various extensions of GCV are proposed for generalized linear models [36], [37], [38]. For logistic regression we use the approach by O'Sullivan et al. [36]: we iterate the Newton-Raphson algorithm until convergence, then, at the final iteration we compute the GCV value on the quadratic approximation (Eq. (22)).

Success Rates. Performance of the templates is evaluated using success rates. The success rate of a template is the percentage of images in which the target object was successfully localized. In both optic nerve head (Subsec. 4.2) and fovea (Subsec. 4.3) detection experiments, a successful detection is defined as such if the detected location $\mathbf{x}^*$ (Eq. (3)) lies within one optic disk radius distance to the actual location. For pupil detection both the left and right eye need to be detected and we therefore use the following normalized error metric

$$e = \frac{\max(d_{left}, d_{right})}{w}, \tag{39}$$

in which w is the (ground truth) distance between the left and right eye, and d_{left} and d_{right} are respectively the distances of detection locations to the left and right eye.

k-Fold Cross Validation. For correct unbiased evaluation, none of the test images are used for training of the templates, nor are they used for parameter optimization. We perform k-fold cross validation: The complete dataset is randomly partitioned into k subsets. Training (patch extraction, parameter optimization and template construction) is done using the data from $k-1$ subsets. Template matching is then performed on the remaining subset. This is done for all k configurations with $k-1$ training subsets and one test subset, allowing us to compute the average performance (success rate) and standard deviation. We set $k = 5$.

4.2 Optic Nerve Head Detection in Retinal Images

Our first application to retinal images is optic nerve head detection. The ONH is one of the key anatomical landmarks in the retina, and its location is often used as a reference point to define regions of interest for the analysis of the retina. The detection hereof is therefore an essential step in many automated retinal image analysis pipelines.

The ONH has two main characteristics: 1) it often appears as a bright disk-like structure on color fundus (CF) images (dark on SLO images), and 2) it is the place from which blood vessels leave the retina. Traditionally, methods have mainly focused on the first characteristic [39], [40], [41]. However, in case of bad contrast of the optic disk, or in the presence of pathology (especially bright lesions, see e.g. Fig. 3), these methods typically fail. Most of the recent ONH detection methods therefore also include the vessel patterns in the analysis; either via explicit vessel segmentation [42], [43], vessel density measures [44], [45], or via additional orientation pattern matching steps [46]. In our method, both the appearance and vessel characteristics are addressed in an efficient integrated template matching approach, resulting in state-of-the-art performance both in terms of success rates and computation time. We target the first characteristic with template matching on $\mathbb{R}^2$. The second is targeted with template matching on $SE(2)$.

4.2.1 Processing Pipeline & Data

Processing Pipeline. First, the images are rescaled to a working resolution of 40 $\mu m/pix$. In our experiments the average resolution per dataset was determined using the average optic disk diameter (which is on average $1.84mm$). The images are normalized for contrast and illumination variations using the method from [47]. Finally, in order to put more emphasis on contextual/shape information, rather than pixel intensities, we apply a soft binarization to the locally normalized (cf. Eq. (31) in Ch. 3 of the supplementary materials) image f via the mapping $\mathrm{erf}(8f)$.

For the orientation score transform we use $N_\theta = 12$ uniformly sampled orientations from 0 to π and lift the image using cake wavelets [6], [7]. For phase-invariant, nonlinear, left-invariant [10], and contractive [48] processing on SE(2), we work with the modulus of the complex valued orientation scores rather than with the complex-valued scores themselves (taking the modulus of quadrature filter responses is an effective technique for line detection, see e.g. Freeman et al. [49]).

Due to differences in image characteristics, training and matching is done separately for the SLO and the color fundus images. For SLO images we use the near infrared channel, for RGB fundus images we use the green channel.

Positive training samples f_i are defined as $N_x \times N_y$ patches, with $N_x = N_y = 251$, centered around true ONH

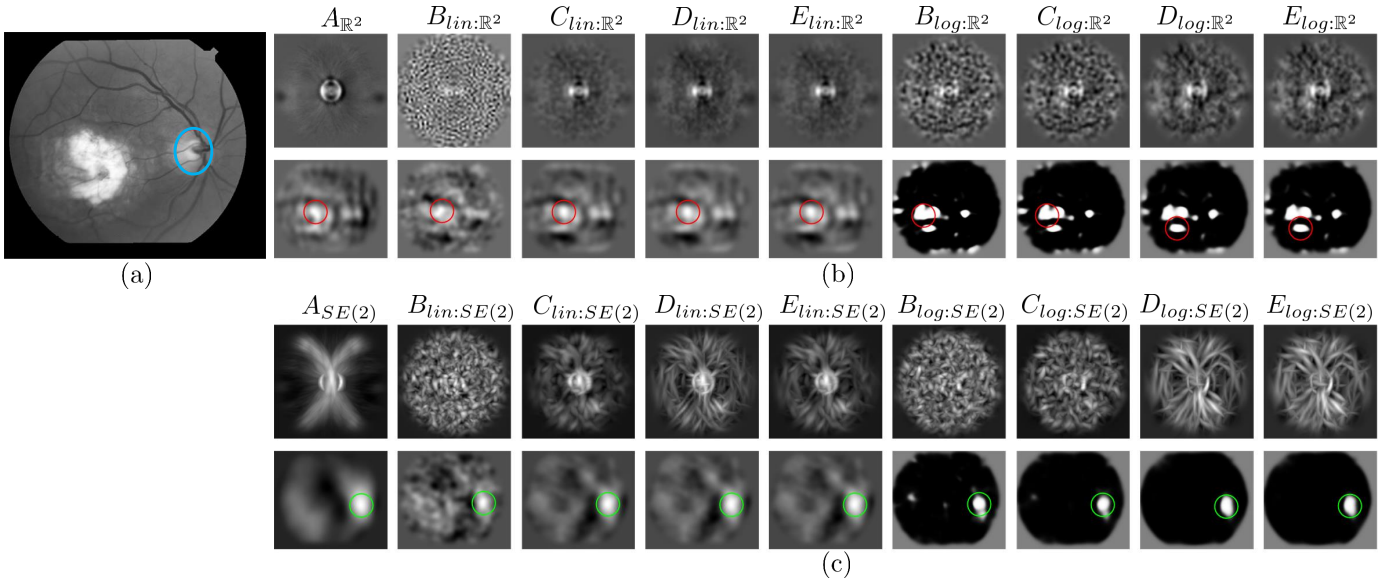

Fig. 3. Overview of trained templates for ONH detection, and their responses to a challenging retinal image. **(a)** The example input image with true ONH location in blue. **(b)** The $\mathbb{R}^2$-type templates (top row) and their responses to the input image (bottom row). **(c)** The maximum intensity projections (over θ) of the $SE(2)$-type templates (top row) and their responses to the input image (bottom row). Detected ONH locations are indicated with colored circles (green = correct, red = incorrect).

TABLE 1
Average template matching results ($\pm$ standard deviation) for optic nerve head detection in 5-fold cross validation, number of failed detections in parentheses.

Template ID	ES (SLO) 208	TC 208	MESSIDOR 1200	DRIVE 40	STARE 81	All Images 1737
			$\mathbb{R}^2$ templates			
$A_{\mathbb{R}^2}$	100.0% ± 0.00% (0)	99.49% ± 1.15% (1)	98.83% ± 0.56% (14)	96.36% ± 4.98% (2)	74.94% ± 9.42% (20)	97.87% ± 0.52% (37)
$B_{lin:\mathbb{R}^2}$	99.09% ± 2.03% (2)	20.35% ± 5.99% (165)	9.67% ± 2.69% (1084)	9.09% ± 12.86%(35)	3.56% ± 3.28% (78)	21.48% ± 2.16% (1364)
$C_{lin:\mathbb{R}^2}$	99.55% ± 1.02% (1)	99.57% ± 0.97% (1)	98.33% ± 0.41% (20)	94.55% ± 8.13% (3)	66.96% ± 16.65%(26)	97.07% ± 0.76% (51)
$D_{lin:\mathbb{R}^2}$	99.55% ± 1.02% (1)	99.57% ± 0.97% (1)	98.42% ± 0.45% (19)	96.36% ± 4.98% (2)	67.53% ± 17.80%(25)	97.24% ± 0.72% (48)
$E_{lin:\mathbb{R}^2}$	99.55% ± 1.02% (1)	99.57% ± 0.97% (1)	98.33% ± 0.29% (20)	96.36% ± 4.98% (2)	66.90% ± 19.25%(26)	97.12% ± 0.84% (50)
$B_{log:\mathbb{R}^2}$	4.36% ± 3.21% (199)	4.59% ± 6.41% (199)	3.17% ± 0.86% (1162)	1.82% ± 4.07% (39)	3.64% ± 8.13% (79)	3.40% ± 0.74% (1678)
$C_{log:\mathbb{R}^2}$	68.69% ± 6.24% (65)	98.10% ± 2.00% (4)	97.75% ± 1.01% (27)	96.36% ± 4.98% (2)	66.94% ± 16.43%(28)	92.74% ± 0.65% (126)
$D_{log:\mathbb{R}^2}$	41.87% ± 6.81% (121)	97.60% ± 1.82% (5)	96.00% ± 1.59% (48)	91.01% ± 8.46% (4)	65.30% ± 10.05%(28)	88.14% ± 1.21% (206)
$E_{log:\mathbb{R}^2}$	58.68% ± 4.48% (86)	97.59% ± 2.48% (5)	97.33% ± 0.96% (32)	93.51% ± 9.00% (3)	67.88% ± 12.61%(27)	91.20% ± 0.95% (153)
			$SE(2)$ templates			
$A_{SE(2)}$	98.57% ± 2.16% (3)	98.95% ± 2.35% (2)	99.58% ± 0.30% (5)	98.18% ± 4.07% (1)	94.22% ± 9.64% (5)	99.08% ± 0.75% (16)
$B_{lin:SE(2)}$	99.06% ± 1.29% (2)	94.75% ± 2.48% (11)	93.74% ± 1.80% (75)	92.05% ± 7.95% (4)	85.63% ± 10.97%(12)	94.01% ± 0.89% (104)
$C_{lin:SE(2)}$	99.06% ± 1.29% (2)	100.0% ± 0.00% (0)	100.0% ± 0.00% (0)	97.50% ± 5.59% (1)	94.00% ± 6.17% (5)	99.54% ± 0.39% (8)
$D_{lin:SE(2)}$	98.60% ± 2.05% (3)	100.0% ± 0.00% (0)	99.67% ± 0.46% (4)	100.0% ± 0.00% (0)	94.00% ± 6.17% (5)	99.31% ± 0.44% (12)
$E_{lin:SE(2)}$	98.60% ± 2.05% (3)	100.0% ± 0.00% (0)	99.67% ± 0.46% (4)	97.50% ± 5.59% (1)	95.11% ± 5.48% (4)	99.31% ± 0.33% (12)
$B_{log:SE(2)}$	87.06% ± 4.20% (27)	77.68% ± 5.36% (46)	84.17% ± 2.25% (190)	80.19% ± 14.87% (9)	75.10% ± 9.81% (21)	83.14% ± 1.78% (293)
$C_{log:SE(2)}$	97.66% ± 2.79% (5)	99.52% ± 1.06% (1)	99.58% ± 0.42% (5)	98.18% ± 4.07% (1)	95.33% ± 7.30% (4)	99.08% ± 0.13% (16)
$D_{log:SE(2)}$	95.22% ± 3.78% (10)	98.50% ± 2.27% (3)	99.25% ± 0.19% (9)	98.18% ± 4.07% (1)	95.33% ± 4.74% (4)	98.45% ± 0.38% (27)
$E_{log:SE(2)}$	97.14% ± 2.61% (6)	99.52% ± 1.06% (1)	99.50% ± 0.35% (6)	98.18% ± 4.07% (1)	94.22% ± 6.82% (5)	98.90% ± 0.48% (19)
			Template combinations (sorted on performance)			
$A_{\mathbb{R}^2}$ + $C_{log:SE(2)}$	100.0% ± 0.00% (0)	100.0% ± 0.00% (0)	99.92% ± 0.19% (1)	98.18% ± 4.07% (1)	98.67% ± 2.98% (1)	99.83% ± 0.26% (3)
$A_{\mathbb{R}^2}$ + $E_{log:SE(2)}$	100.0% ± 0.00% (0)	100.0% ± 0.00% (0)	99.83% ± 0.23% (2)	98.18% ± 4.07% (1)	98.67% ± 2.98% (1)	99.77% ± 0.24% (4)
$A_{\mathbb{R}^2}$ + $D_{log:SE(2)}$	100.0% ± 0.00% (0)	100.0% ± 0.00% (0)	99.83% ± 0.23% (2)	98.18% ± 4.07% (1)	98.67% ± 2.98% (1)	99.77% ± 0.24% (4)
$C_{lin:SE(2)}$ + $E_{log:SE(2)}$	99.55% ± 1.02% (1)	100.0% ± 0.00% (0)	99.83% ± 0.23% (2)	100.0% ± 0.00% (0)	96.44% ± 3.28% (3)	99.65% ± 0.13% (6)
$C_{lin:SE(2)}$ + $C_{log:SE(2)}$	99.55% ± 1.02% (1)	100.0% ± 0.00% (0)	99.92% ± 0.19% (1)	98.18% ± 4.07% (1)	96.44% ± 3.28% (3)	99.65% ± 0.13% (6)
...			...			...
*$A_{SE(2)}$ + $C_{lin:SE(2)}$	99.55% ± 1.02% (1)	100.0% ± 0.00% (0)	100.0% ± 0.00% (0)	98.18% ± 4.07% (1)	94.22% ± 6.82% (5)	99.60% ± 0.26% (7)
...			...			...
†$A_{\mathbb{R}^2}$ + $A_{SE(2)}$	100.0% ± 0.00% (0)	100.0% ± 0.00% (0)	99.66% ± 0.35% (4)	98.18% ± 4.07% (1)	88.42% ± 11.23% (9)	99.19% ± 0.63% (14)
...			...			...

**Best template combination that does not rely on logistic regression.* †*Best template combination that does not rely on template optimization.*

location in each image. For every image, a negative sample is defined as an image patch centered around random location in the image that does not lie within one optic disk radius distance to the true ONH location. An exemplary ONH patch is given in Fig. 1. For the B-spline expansion of the templates we set $N_k = N_l = 51$ and $N_m = 12$.

Data. In our experiments we made use of both publicly available data, and a private database. The private database consists of 208 SLO images taken with an EasyScan fundus camera (i-Optics B.V., the Netherlands) and 208 CF images taken with a Topcon NW200 (Topcon Corp., Japan). Both cameras were used to image both eyes of the same patient, taking an ONH centered image, and a fovea centered image per eye. The two sets of images are labeled as "ES" and "TC" respectively. The following (widely used) public databases are also used: MESSIDOR (http://messidor.crihan.fr/index-en.php), DRIVE (http://www.isi.uu.nl/Research/Databases/DRIVE) and STARE (http://www.ces.clemson.edu/~ahoover/stare), consisting of 1200, 40 and 81 images respectively. For each image, the circumference of the ONH was annotated, and parameterized by an ellipse. The annotations for the MESSIDOR dataset were kindly made available by the authors of [50] (http://www.uhu.es/retinopathy). The ONH contour in the remaining images were manually outlined by ourselves. The annotations are made available on our website. The images in the databases contain a mix of good quality healthy images, and challenging diabetic retinopathy cases. Especially MESSIDOR and STARE contain challenging images.

4.2.2 Results and Discussion

The templates. The different templates for ONH detection are visualized in Fig. 3. The $SE(2)$ templates are visualized using maximum intensity projections over θ. In this figure we have also shown template responses to an example image. Visually one can clearly recognize the typical disk shape in the $\mathbb{R}^2$ templates, whereas the $SE(2)$ templates also seem to capture the typical pattern of outward radiating blood vessels (compare e.g. $A_{\mathbb{R}^2}$ with $A_{SE(2)}$). Indeed, when applied to a retinal image, where we took an example with an optic disk like pathology, we see that the $\mathbb{R}^2$ templates respond well to the disk shape, but also (more strongly) to the pathology. In contrast, the $SE(2)$ templates respond mainly to vessel pattern and ignore the pathology. We also see, as expected, a smoothing effect of gradient based regularization (D and E) in comparison to standard $\mathbb{L}_2$-norm regularization (C) and no regularization (B). Finally, in comparison to linear regression templates, the logistic regression templates have a more binary response due to the logistic sigmoid mapping.

Detection results. Table 1 gives a breakdown of the quantitative results for the different databases used in the experiments. The templates are grouped in $\mathbb{R}^2$ templates, $SE(2)$ templates, and combination of templates. Within these groups, they are further divided in average, linear regression, and logistic regression templates. The best overall performance within each group is highlighted in gray.

Overall, we see that the $SE(2)$ templates out-perform their $\mathbb{R}^2$ equivalents, and that combinations of the two types of templates give best results. The two types are nicely complementary to each other due to disk-like sensitivity of the $\mathbb{R}^2$ templates and the vessel pattern sensitivity of the $SE(2)$ templates. If one of the two ONH characteristics is less obvious (as is e.g. for the disk-shape in Fig. 3), the other can still be detected. Also, the failures of $\mathbb{R}^2$ templates are mainly due to either distracting pathologies in the retina, or poor contrast of the optic disk. As reflected by the increased performance of $SE(2)$ templates over $\mathbb{R}^2$ templates, a more stable pattern seems to be the vessel pattern.

From Table 1 we also deduce that the individual performances of the linear regression templates outperform the logistic regression templates. Moreover, the average templates give best individual performance, which indicates that with our effective template matching framework good performance can already be achieved with basic templates. However, we also see that low performing individual templates can prove useful when combining templates. In fact, we see that combinations with all linear $\mathbb{R}^2$ templates are highly ranked, and for the $SE(2)$ templates it is mainly the logistic regression templates. This can be explained by the binary nature of the logistic templates: even when the maximum response of the templates is at an incorrect location, the difference with the correct location is often small. The $\mathbb{R}^2$ template then adds to the sensitivity and precision. The best results obtained with untrained templates was a 99.19% success rate (14 fails), and with the overall best template combination we obtained a 99.83% success rate (3 fails).

State of the art. In Table 2 we compare our results on the publicly available benchmark databases MESSIDOR, DRIVE and STARE, with the most recent methods for ONH detection (sorted from oldest to newest from top to bottom). In this comparison, our best performing method ($A_{\mathbb{R}^2}+C_{log:SE(2)}$) performs better than or equally well as the best methods from literature. We have also listed the computation times, and see that our method is also ranked as one of the fastest methods for ONH detection. The average computation time, using our experimental implementation in Wolfram *Mathematica* 10.4, was 0.5 seconds per image on a computer with an Intel Core i703612QM CPU and 8GB memory. A full breakdown of timings of the processing pipeline is given in the supplementary materials Sec. 4.

TABLE 2
Comparison to state of the art: Optic nerve head detection success rates, the number of fails (in parentheses), and computation times.

Method	MESSIDOR	DRIVE	STARE	Time (s)
Lu [51]	99.8% (3)		98.8% (1)	5.0
Lu et al. [39]		97.5% (1)	96.3% (3)	40.0
Yu et al. [44]	99.1% (11)			4.7
Aquino et al. [40]	99.8% (14)			1.7
Giachetti et al. [45]	99.7% (4)			5.0
Ramakanth et al. [27]	99.4% (7)	100% (0)	93.83% (5)	0.2
Marin et al. [43]	99.8% (3)			5.4†
Dashtbozorg et al. [41]	99.8% (3)			10.6†
Proposed	99.9% (1)	97.8% (1)	98.8% (1)	0.5

†*Timings include simultaneous disk segmentation.*

4.3 Fovea Detection in Retinal Images

Our second application to retinal images is for the detection of the fovea. The fovea is the location in the retina which is responsible for sharp central vision. It is characterized by a

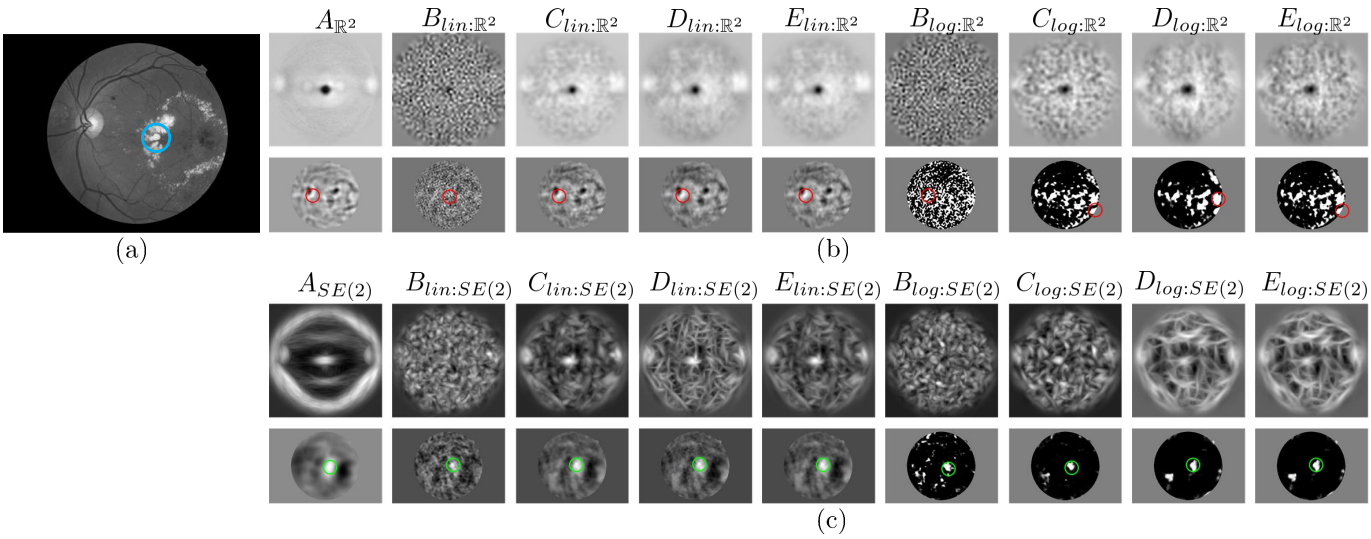

Fig. 4. Overview of trained templates for fovea detection, and their responses to a challenging retinal image. **(a)** The example input image with true fovea location in blue. **(b)** The $\mathbb{R}^2$-type templates (top row) and their responses to the input image (bottom row). **(c)** The maximum intensity projections (over θ) of the $SE(2)$-type templates (top row) and their responses to the input image (bottom row). Detected fovea locations are indicated with colored circles (green = correct, red = incorrect).

small depression in thickness of the retina, and on healthy retinal images it often appears as a darkened area. Since the foveal area is responsible for detailed vision, this area is weighted most heavily in grading schemes that describe the severity of a disease. Therefore, correct localization of the fovea is essential in automatic grading systems [52].

Methods for the detection of the fovea heavily rely on contextual features in the retina [28], [45], [53], [54], [55], and take into account the prior knowledge that 1) the fovea is located approximately 2.5 optic disk diameters lateral to the ONH center, that 2) it lies within an avascular zone, and that 3) it is surrounded by the main vessel arcades. All of these methods restrict their search region for the fovea location to a region relative to the (automatically detected) ONH location. To the best of our knowledge, the proposed detection pipeline is the first that is completely independent of vessel segmentations and ONH detection. This is made possible due to the fact that anatomical reference patterns, in particular the vessel structures, are generically incorporated in the learned templates via data representations in orientation scores.

4.3.1 Processing Pipeline & Data

Processing Pipeline. The proposed fovea detection pipeline is the same as for ONH detection, however, now the positive training samples f_i are centered around the fovea.

Data. The proposed fovea detection method is validated on our (annotated) databases "ES" and "TC", each consisting of 208 SLO and 208 color fundus images respectively (cf. Subsec.4.2.1). We further test our method on the most used publicly available benchmark dataset MESSIDOR (1200 images). Success rates were computed based on the fovea annotations kindly made available by the authors of [28].

4.3.2 Results and Discussion

The templates. Akin to Fig. 3, in Fig. 4 the trained fovea templates and their responses to an input image are visualized. The $\mathbb{R}^2$ templates seem to be more tuned towards the dark (isotropic) blob like appearance of the fovea, whereas in the $SE(2)$ templates one can also recognize the pattern of vessels surrounding the fovea (compare $A_{\mathbb{R}^2}$ with $A_{SE(2)}$). To illustrate the difference between these type of templates, we selected an image in which the fovea location is occluded with bright lesions. In this case the method has to rely on contextual information (e.g. the blood vessels). Indeed, we see that the $\mathbb{R}^2$ templates fail due to the absence of a clear foveal blob shape, and that the $SE(2)$ templates correctly identify the fovea location. The effect of regularization is also clearly visible; no regularization (B) results in noisy templates, standard $\mathbb{L}_2$ regularization (C) results in more stable templates, and smoothed regularization (D and E) results in smooth templates. In templates $D_{SE(2)}$ and $E_{SE(2)}$ we see that more emphasis is put on line structures.

Detection results. A full overview of individual and combined template performance is discussed in the supplementary materials, here we only provide a summary. Again there is an improvement using $SE(2)$ templates over $\mathbb{R}^2$ templates, although the difference is smaller than in the ONH application. Apparently both the dark blob-like appearance ($\mathbb{R}^2$ templates) and vessel patterns ($SE(2)$ templates) are equally reliable features of the fovea. A combination of templates leads to improved results and we conclude that the templates are again complementary to each other. Furthermore, again linear regression performs better than logistic regression. In fovea detection we do observe a large improvement of template training over basic averaging: 1529 of 1616 (94.6%) successful detections with $C_{lin:SE(2)}$ versus 1488 (92.1%) with $A_{SE(2)}$. The best performing $\mathbb{R}^2$ template was $A_{\mathbb{R}^2}$ (65.6%), the best $SE(2)$ template was $C_{lin:SE(2)}$ (94.6%). The best combination of templates was $C_{lin:\mathbb{R}^2} + C_{log:SE(2)}$ with 1605 (99.3%) detections. When using non-optimized templates 1588 (98.3%) successful detections were achieved (with $A_{\mathbb{R}^2} + A_{SE(2)}$).

State of the art. In Table 3 we compared our results on

TABLE 3
Comparison to state of the art: Fovea detection success rates, the number of fails (in parentheses), and computation times.

Method	MESSIDOR	Time (s)
Niemeijer et al. [28], [55]	97.9% (25)	7.6†
Yu et al. [54]	95.0%* (60)	3.9†
Gegundez-Arias et al. [28]	96.9% (37)	0.9
Giachetti et al. [45]	99.1% (11)	5.0†
Aquino [53]	98.2% (21)	10.9†
Proposed	99.7% (3)	0.5

*Success-criterion based on half optic radius.
†Timing includes ONH detection.

the publicly available benchmark database MESSIDOR with the most recent methods for fovea detection (sorted from oldest to newest from top to bottom). In this comparison, our best performing method ($C_{lin:\mathbb{R}^2} + C_{log:SE(2)}$) quite significantly outperforms the best methods from literature. Furthermore, our detection pipeline is also the most efficient one; the computation time for fovea detection is the same as for ONH detection, which is 0.5 seconds.

4.4 Pupil Detection

Our third application is that of pupil localization in regular camera images, which is relevant in many applications as they provide important visual cues for face detection, face recognition, and understanding of facial expressions. In particular in gaze estimation the accurate localization of the pupil is essential. Eye detection and tracking is however challenging due to, amongst others: occlusion by the eyelids and variability in size, shape, reflectivity or head pose.

Many pupil localization algorithms are designed to work on periocular images, these are close-up views of the eyes. Such images can be acquired by dedicated eye imaging devices, or by means of cropping a full facial image (see Fig. 5(a)). We will consider both the problem of detection pupils in periocular images and the more difficult problem of detection in full images.

We compare our method against the seven most recent pupil detection methods from literature, for a full overview see [56] and [57]. A method similar to our $\mathbb{R}^2$ approach in the sense that it is also based on 2D linear filtering is the method by Kroon et al. [58]. In their paper templates are obtained via linear discriminant analysis of pupil images. Asteriada et al. [59] detect the pupil by matching templates using features that are based on distances to the nearest strong (facial) edges in the image. Campadelli et al. [60] use a supervised approach with a SVM classifier and Haar wavelet features. The method by Timm et al. [61] is based on searching for gradient fields with a circular symmetry. Valenti et al. [62] use a similar approach but additionally include information of isophote curvature, with supervised refinement. Markus et al. [57] employ a supervised approach using an ensemble of randomized regression trees. Leo et al. [56] employ a completely unsupervised approach similar to those in [61], [62], but additionally include analysis of self-similarity.

A relevant remark is that all of the above mentioned methods rely on prior face detection, and restrict their search region to periocular images. Our method works completely stand alone, and can be used on full images.

4.4.1 Processing Pipeline & Data

Processing Pipeline. Interestingly, we could again employ the same processing pipeline (including local normalization via [47]) which was used for ONH and fovea detection. In our experiments we train templates for the left and right eye separately.

Data. We validated our pupil detection approach on the publicly available BioID database (http://www.bioid.com), which is generally considered as one of the most challenging and realistic databases for pupil detection in facial images. The database consists of 1521 frontal face grayscale images with significant variation in illumination, scale and pose.

4.4.2 Results and Discussion

The templates. Fig. 5(b) and (c) show respectively the trained $\mathbb{R}^2$ and $SE(2)$ templates for pupil detection of the right eye, and their filtering response to the input image in Fig. 5(a). Here the trained $\mathbb{R}^2$ templates seemed to capture the pupil as a small blob in the center of the template, but apart from that no real structure can be observed. In the average template we do however clearly see structure in the form of an "average face". The $SE(2)$ templates reveal structures that resemble the eyelids in nearly all templates. The linear regression templates look sharper and seem to contain more detail than the average template, and the logistic regression templates seem to take a good compromise between course features and details.

Detection results. We again refer to the supplementary materials for a full benchmarking analysis, in summary we observed the following. In terms of success rates we see a similar pattern as with the ONH and fovea application, however, here we see that the learned templates (C,D and E) significantly outperform the average templates, and that logistic regression leads to better templates than using linear regression (94.0% success rate for $C_{log:SE(2)}$ vs 87.2% for $D_{lin:SE(2)}$). Overall, the $SE(2)$ templates outperform the $\mathbb{R}^2$ templates, linear regression templates outperform the average template, and logistic regression templates outperform linear regression templates. The best $\mathbb{R}^2$ template was $D_{lin:\mathbb{R}^2}$ with 1151 of 1521 detections (75.7%), the best $SE(2)$ template was $C_{log:SE(2)}$ (94.0%). The best combination of templates was $D_{lin:\mathbb{R}^2}$ with $E_{lin:SE(2)}$ (95.6%). Without template training (i.e., using average templates A) the performance was only 68.2%. Success rates using the best template combination are given in Fig. 5(d) and (e). The processing time for detecting both pupils simultaneously was on average 0.4 seconds per image.

State of the art. In Fig. 5(d) we compared our approach to the two most recent pupil detection methods from literature for several normalized error thresholds. Here we see that with allowed errors of 0.1 (blue circles Fig. 5(a)) and higher our method competes very well with the state of the art, despite the fact that our generic method is not adapted to the application. Further application specific tuning and preprocessing could be applied to improve precision (for $e \ll 0.1$), but this is beyond the scope of this article. Moreover, we see that our method can be used on full images instead of the periocular images without much loss in performance. The fact that our method is still very accurate on full image processing shows that it can be used as a preprocessing step for other applications.

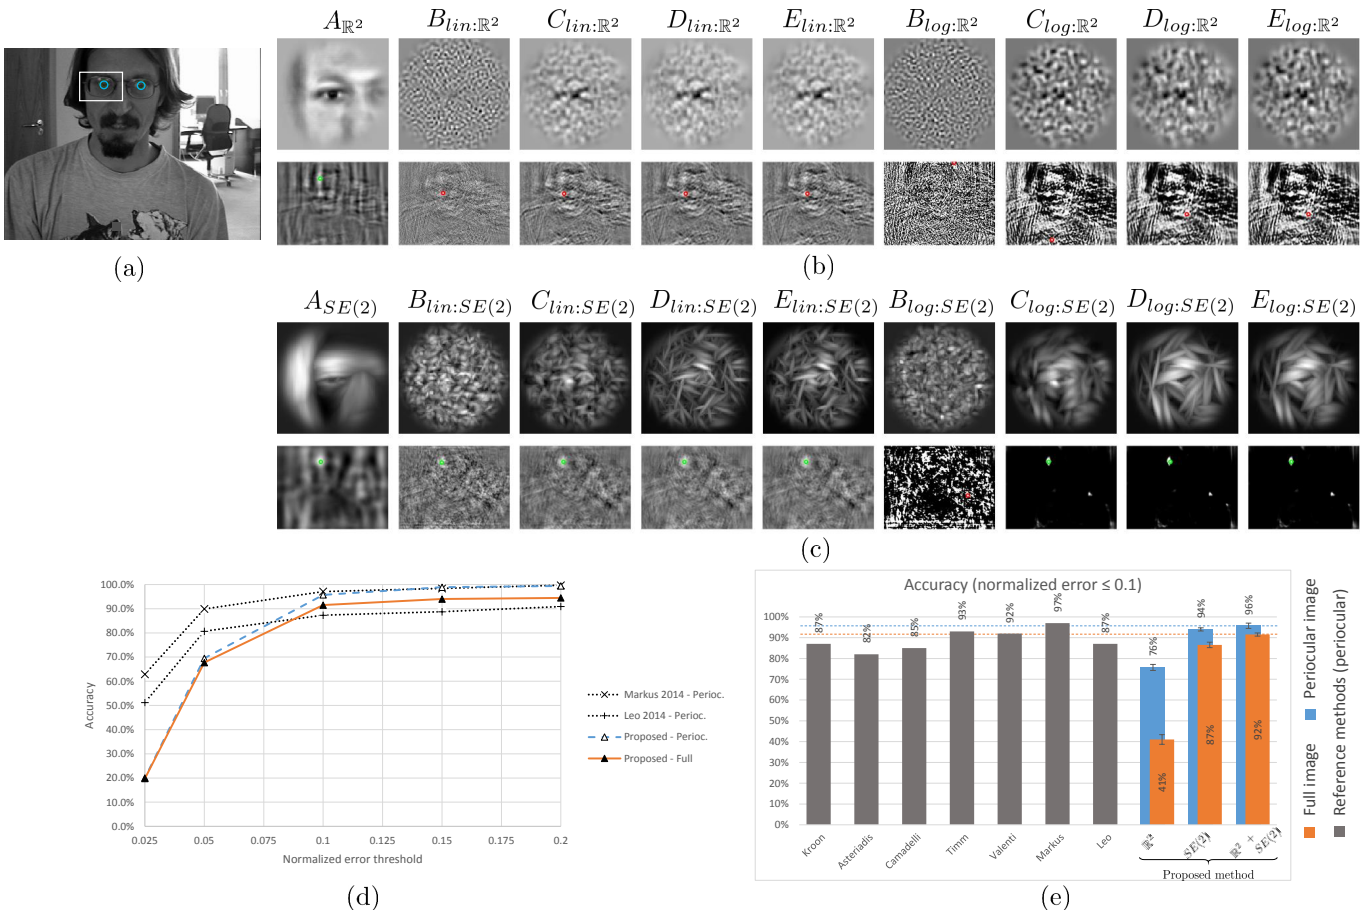

Fig. 5. Overview of trained templates for right-eye pupil detection, and their responses to a challenging image from the BioID database. **(a)** The example input image with true pupil locations (blue circle with a radius that corresponds to a normalized error threshold of 0.1, see Eq. 39. The white square indicates the periocular image region for the right eye. **(b)** The $\mathbb{R}^2$-type templates (top row) and their responses to the input image (bottom row). **(c)** The maximum intensity projections (over θ) of the $SE(2)$-type templates (top row) and their responses to the input image (bottom row). Detected pupil locations are indicated with colored circles (green = correct, red = incorrect, based on a normalized error threshold of 0.1). **(d)** Accuracy curves generated by varying thresholds on the normalized error, in comparison with the two most recent methods from literature. **(e)** Accuracy (at a normalized error threshold of 0.1) comparison with pupil detection methods from literature.

If Fig. 5(e) we compared our approach to the seven most recent methods from literature (sorted from old to new). Here we see that the only method outperforming our method, at standard accuracy requirements ($e \leq 0.1$), is the method by Markus et al. [57]. Even when considering processing of the full images the only other method that outperforms ours is the method by Timm et al. [61], whose performance is measured using periocular images.

4.5 General Observations

The application of our method to the three problems (ONH, fovea and pupil detection) showed the following:

1) State-of-the-art performance was achieved on three different applications, using a single (generic) detection framework and without application specific parameter adaptations.
2) Cross correlation based template matching via data representations on $SE(2)$ improves results over standard $\mathbb{R}^2$ filtering.
3) Trained templates, obtained using energy functionals of the form (1), often perform better than basic average templates. In particular in pupil detection the optimization of templates proved to be essential.
4) Our newly introduced logistic regression approach leads to improved results in pupil detection via single templates. When combining templates we observe only a small improvement of choosing logistic regression (instead of linear regression) for the application of ONH and fovea detection.
5) Regularization in both linear and logistic regression is important. Here both ridge and smoothing regularization priors have complementary benefits.
6) Our method does not rely on any other detection systems (such as ONH detection in the fovea application, or face detection in the pupil detection), and still performs well compared to methods that do.
7) Our method is fast and parallelizable as it is based on inner products, as such it could be efficiently implemented using convolutions.

5 Conclusion

In this paper we have presented an efficient cross-correlation based template matching scheme for the detection of com-

bined orientation and blob patterns. Furthermore, we have provided a generalized regression framework for the construction of templates. The method relies on data representations in orientation scores, which are functions on the Lie group $SE(2)$, and we have provided the tools for proper smoothing priors via resolvent hypo-elliptic diffusion processes on $SE(2)$ (solving time-integrated hypo-elliptic Brownian motions on $SE(2)$). The strength of the method was demonstrated with two applications in retinal image analysis (the detection of the optic nerve head (ONH), and the detection of the fovea) and additional experiments for pupil detection in regular camera images. In the retinal applications we achieved state-of-the-art results with an average detection rate of 99.83% on 1737 images for ONH detection, and 99.32% on 1616 images for fovea detection. Also on pupil detection we obtained state-of-the-art performance with a 95.86% success rate on 1521 images. We showed that the success of the method is due to the inclusion of both intensity and orientation features in template matching. The method is also computationally efficient as it is entirely based on a sequence of convolutions (which can be efficiently done using fast Fourier transforms). These convolutions are parallelizable, which can further speed up our already fast experimental *Mathematica* implementations that are publicly available at http://erikbekkers.bitbucket.org/TMSE2.html. In future work we plan to investigate the applicability of smoothing on $SE(2)$ in variational settings, as this could also be used in (sparse) line enhancement and segmentation problems.

ACKNOWLEDGMENTS

The authors would like to thank the groups that kindly made available the benchmark datasets and annotations. The authors gratefully acknowledge Gonzalo Sanguinetti (TU/e) for fruitful discussions and feedback on this manuscript. The research leading to the results of this article has received funding from the European Research Council under the European Community's 7th Framework Programme (FP7/20072014)/ERC grant agreement No. 335555. This work is also part of the Hé Programme of Innovation Cooperation, which is (partly) financed by the Netherlands Organisation for Scientific Research (NWO).

APPENDIX A
PROBABILISTIC INTERPRETATION OF THE SMOOTHING PRIOR IN $SE(2)$

In this section we relate the $SE(2)$ smoothing prior to time resolvent hypo-elliptic[1] diffusion processes on $SE(2)$. First we aim to familiarize the reader with the concept of resolvent diffusions on $\mathbb{R}^2$ in Subsec. A.1. Then we pose in Subsec. A.2 a new problem (the single patch problem), which is a special case of our $SE(2)$ linear regression problem, that we use to link the left-invariant regularizer to tue resolvents of hypo-elliptic diffusions on $SE(2)$.

A.1 Resolvent Diffusion Processes

A classic approach to noise suppression in images is via diffusion regularizations with PDE's of the form [34]

$$\left\{ \begin{array}{ll} \frac{\partial}{\partial \tau} u & = \Delta u, \\ u|_{\tau=0} & = u_0, \end{array} \right. \tag{40}$$

where Δ denotes the Laplace operator. Solving (40) for any diffusion time $\tau > 0$ gives a smoothed version of the input u_0. The time-resolvent process of the PDE is defined by the Laplace transform with respect to τ; time τ is integrated out using a memoryless negative exponential distribution $P(\mathcal{T} = \tau) = \alpha e^{-\alpha\tau}$. Then, the time integrated solutions

$$t(\mathbf{x}) = \alpha \int_0^\infty u(\mathbf{x}, \tau) e^{-\alpha\tau} \mathrm{d}\tau,$$

with decay parameter α, are in fact the solutions

$$t = \underset{t \in \mathbb{L}_2(\mathbb{R}^2)}{\operatorname{argmin}} \left[\|t - t_0\|^2_{\mathbb{L}_2(\mathbb{R}^2)} + \lambda \int_{\mathbb{R}^2} \|\nabla t(\tilde{\mathbf{x}})\|^2 \, \mathrm{d}\tilde{\mathbf{x}} \right], \tag{41}$$

with $\lambda = \alpha^{-1}$, and corresponding Euler-Lagrange equation

$$(I - \lambda\Delta)t = t_0 \quad \Leftrightarrow \quad t = \lambda^{-1} \left(\frac{1}{\lambda} - \Delta \right)^{-1} t_0, \tag{42}$$

to which we refer as the "resolvent" equation [63], as it involves operator $(\alpha I - \Delta)^{-1}$, $\alpha = \lambda^{-1}$. In the next subsections, we follow a similar procedure with $SE(2)$ instead of $\mathbb{R}^2$, and show how the smoothing regularizer in Eq. (28) and (30) of the main article relates to Laplace transforms of hypo-elliptic diffusions on the group $SE(2)$ [9], [10].

A.2 The Fundamental Single Patch Problem

In order to grasp what the (anisotropic regularization term) in Eq. (28) and (30) of the main article actually means in terms of stochastic interpretation/probabilistic line propagation, let us consider the following single patch problem and optimize

$$\begin{aligned} \mathcal{E}_{sp}(T) &= \left| \left(G_s *_{\mathbb{R}^2} T(\cdot, \cdot, \theta_0) \right)(\mathbf{x}_0) - 1 \right|^2 \\ &+ \lambda \int_{\mathbb{R}^2} \int_0^{2\pi} \|\nabla T(\tilde{\mathbf{x}}, \tilde{\theta})\|^2_D \mathrm{d}\tilde{\mathbf{x}} \mathrm{d}\tilde{\theta} + \mu \|T\|^2_{\mathbb{L}_2(SE(2))}, \end{aligned} \tag{43}$$

with $(\mathbf{x}_0, \theta_0) = g_0 := (x_0, y_0, \theta_0) \in SE(2)$ the fixed center of the template, and with spatial Gaussian kernel

$$G_s(\mathbf{x}) = \frac{1}{4\pi s} e^{-\frac{\|\mathbf{x}\|^2}{4s}}.$$

Regarding this problem, we note the following:

- In the original problem (28) of the main article we take $N = 1$, with

$$U_{f_1}(x, y, \theta) = G_s(x - x_0, y - y_0) \, \delta_{\theta_0}(\theta) \tag{44}$$

  representing a local spatially smoothed spike in $SE(2)$, and set $y_1 = 1$. The general single patch case (for arbitrary U_{f_1}) can be deduced by superposition of such impulse responses.
- We use $\mu > 0$ to suppress the output elsewhere.
- We use $0 < s \ll 1$. This minimum scale due to sampling removes the singularity at $(\mathbf{0}, 0)$ from the kernel that solves (43), as proven in [9].

Theorem 1. *The solution to the single patch problem (43) coincides up to scalar multiplication with the time integrated hypo-elliptic Brownian motion kernel on $SE(2)$ (depicted in Fig. 6).*

Proof. We optimize $\mathcal{E}_{sp}(T)$ over the set $\mathcal{S}(SE(2))$ of all functions $T : SE(2) \to \mathbb{R}$ that are bounded and on $SE(2)$, infinitely differentiable on $SE(2) \setminus \{g_0\}$, and rapidly decreasing in spatial direction, and 2π periodic in θ. We omit topological details on function spaces and Hörmanders condition [64]. Instead, we directly proceed with applying the Euler-Lagrange technique to the single patch problem:

$$\begin{aligned} \forall_{\delta \in \mathcal{S}(SE(2))} : \lim_{\epsilon \downarrow 0} \left\{ \frac{\mathcal{E}_{sp}(T + \epsilon\delta) - \mathcal{E}_{sp}(T)}{\epsilon} \right\} = 0 \Leftrightarrow \\ (S_s^* S_s + \lambda R + \mu I)T = S_s^* y_1 = S_s^* 1, \end{aligned} \tag{45}$$

with linear functional (distribution) $\mathcal{S}_s$ given by

$$(S_s T) = (G_s *_{\mathbb{R}^2} T(\cdot, \theta_0))(\mathbf{x}_0),$$

and with regularization operator R given by

$$R = -\Delta_{SE(2)} := -(D_{\theta\theta}\partial_\theta^2 + D_{\xi\xi}\partial_\xi^2 + D_{\eta\eta}\partial_\eta^2) \geq 0.$$

Note that $\lim_{s \to 0} S_s = \delta_{(\mathbf{x}_0, \theta_0)}$ in distributional sense, and that the constraint $s > 0$ is crucial for solutions T to be bounded at $(\mathbf{x}_0, \theta_0)$. By definition the adjoint operator S_s^* is given by

$$\begin{aligned} (S_s^* y, T)_{\mathbb{L}_2(SE(2))} &= (y, S_s T) = y \textstyle\int_{\mathbb{R}^2} G_s(\mathbf{x} - \mathbf{x}_0) T(\mathbf{x}, \theta_0) \, \mathrm{d}\mathbf{x} \\ &= y \int_0^{2\pi} \int_{\mathbb{R}^2} G_s(\mathbf{x} - \mathbf{x}_0) \delta_{\theta_0}(\theta) T(\mathbf{x}, \theta) \, \mathrm{d}\mathbf{x} \mathrm{d}\theta, \\ &= (y \, G_s(\cdot - \mathbf{x}_0) \delta_{\theta_0}(\cdot), T)_{\mathbb{L}_2(SE(2))} \end{aligned}$$

and thereby we deduce that

$$\begin{aligned} (S_s^* y)(\mathbf{x}, \theta) &= y \, G_s(\mathbf{x} - \mathbf{x}_0) \delta_{\theta_0}(\theta), \\ S_s^*(S_s T) &= T_0^s \, G_s(\mathbf{x} - \mathbf{x}_0) \delta_{\theta_0}(\theta), \end{aligned}$$

with $\infty > T_0^s := (G_s *_{\mathbb{R}^2} T(\cdot, \theta_0))(\mathbf{x}_0) > 1$ for $0 < s \ll 1$. The Euler-Lagrange equation (45) becomes

$$(-\lambda\Delta_{SE(2)} + \mu I)T = (1 - T_0^s) G_s(\mathbf{x} - \mathbf{x}_0) \delta_{\theta_0}(\theta).$$

1. This diffusion process on $SE(2)$ is called *hypo-elliptic* as its generator equals $(\partial_\xi)^2 + D_{\theta\theta}(\partial_\theta)^2$ and diffuses only in 2 directions in a 3D space. This boils down to a sub-Riemannian manifold structure [9], [33]. Smoothing in the missing (∂_η) direction is achieved via the commutator: $[\partial_\theta, \cos\theta\partial_x + \sin\theta\partial_y] = -\sin\theta\partial_x + \cos\theta\partial_y$.

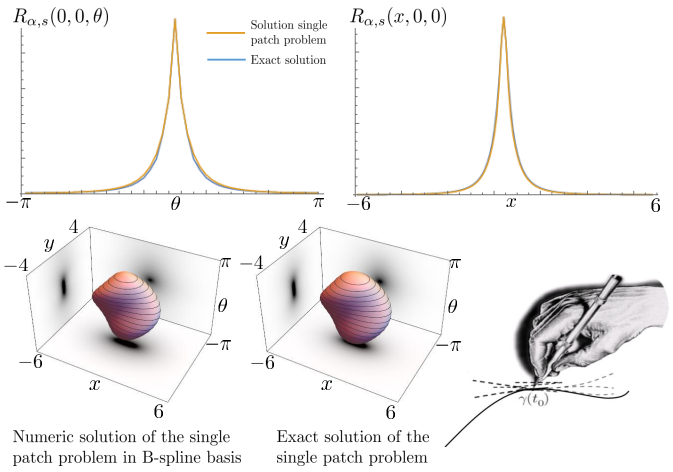
$R_{\alpha,s}(0,0,\theta)$
Solution single patch problem
Exact solution
$R_{\alpha,s}(x,0,0)$
Numeric solution of the single patch problem in B-spline basis
Exact solution of the single patch problem

Fig. 6. Top row: Comparison of kernel $R_{\alpha,s}(x,y,\theta)$ along respectively the θ and x axis. Bottom row: Isosurface of the kernel computed by solving the fundamental single patch problem (43), the exact solution, and an illustration of the drunkman's pencil. For Monte Carlo simulations of the drunkman's pencil see the supplementary materials.

Now, when setting $T_{new} = \frac{T}{1-T^s}$ we arrive at the hypo-elliptic resolvent equation on $SE(2)$:

$$\begin{aligned}(-\lambda\Delta_{SE(2)} + \mu I)T_{new} = (G_s *_{\mathbb{R}^2} \delta_{\mathbf{x}_0})\delta_{\theta_0} \quad &\Leftrightarrow \\ T_{new} = \left(-\lambda\Delta_{SE(2)} + \mu I\right)^{-1} e^{s\Delta_{\mathbb{R}^2}}\delta_{g_0} \quad &(46)\\ = e^{s\Delta_{\mathbb{R}^2}}\left(-\lambda\Delta_{SE(2)} + \mu I\right)^{-1}\delta_{g_0}\end{aligned}$$

where we write $e^{s\Delta_{\mathbb{R}^2}} f = G_s *_{\mathbb{R}^2} f$ for the diffusion operator, to stress the vanishing commutators

$$[e^{s\Delta_{\mathbb{R}^2}}, \Delta_{SE(2)}] = e^{s\Delta_{\mathbb{R}^2}}\Delta_{SE(2)} - \Delta_{SE(2)}e^{s\Delta_{\mathbb{R}^2}} = 0,$$

which directly follows from $[\Delta_{\mathbb{R}^2}, \Delta_{SE(2)}] = 0$. In fact, from these vanishing commutators one can deduce that, thanks to the isotropy of Gaussian kernel, blurring with inner-scale $s > 0$ can be done either before applying the resolvent operator or after (as seen in (46)).

The solutions T_{new} are precisely the probabilistic kernels $R_{\alpha,s} : SE(2) \to \mathbb{R}$ for time integrated contour enhancements studied in [9], [10]. In fact we see that

$$T_{new}(g) = \mu^{-1} R_{\alpha,s}(g_0^{-1}g),$$

where $R_{\alpha,s} = (I - \alpha^{-1}\Delta_{SE(2)})^{-1}e^{s\Delta_{\mathbb{R}^2}}\delta_{(\mathbf{0},0)}$ (i.e., the impuls response of the resolvent operator) denotes the time-integration of the hypo-elliptic diffusion kernel $K_{\tau,s} = e^{\tau\Delta_{SE(2)}}e^{s\Delta_{\mathbb{R}^2}}\delta_{(\mathbf{0},0)}$:

$$R_{\alpha,s}(g) = \alpha\int_0^\infty K_{\tau,s}(g)\, e^{-\alpha\tau}\,\mathrm{d}\tau,$$

for which 3 different exact analytic formulas are derived in [10]. The kernel $R_{\alpha,s}(\mathbf{x},\theta)$ denotes the probability density of finding a random brush stroke (regardless its traveling time) at location $\mathbf{x}$ with orientation θ given that a 'drunkman's pencil' starts at $g = (\mathbf{0},0)$ at time zero. Here the traveling time τ of the random pencil is assumed to be negatively exponentially distributed with expectation α^{-1}. □

A.3 Expansion in B-splines

Now we consider the B-spline expansions (Eq. (34) in the main article) and apply our optimization algorithm (cf. Subsec. 2.4 of the main article) to the single patch problem (43), with $(\mathbf{x}_0,\theta_0) = (\mathbf{0},0)$. Here we no longer need a smoothing with a continuous Gaussian G_s, as expansion in the B-spline basis already includes regularization. Now we set for the smooth spike $U_{f_1}(x,y,\theta) = B^n\left(\frac{x}{s_k}\right)B^n\left(\frac{y}{s_l}\right)B^n\left(\frac{\theta \bmod 2\pi}{s_m}\right)$, and we thus approximate spikes by the same B-spline basis in which we expressed our templates. We accept extra regularization (like we did with the Gaussian in the previous section) and choose to represent a spike by a normal B-spline. After all, via the central limit theorem B-splines converge to Gaussians when increasing n. We also considered to instead use the fundamental B-Spline [20, Fig. 2], which is sharper but suffers from oscillations, yielding less favorable results.

In our normal B-spline setting, this choice of smooth spike representation (cf. Eq. (14) in the main article) leads to the following equations

$$(S^\dagger S + \lambda R + \mu I)T = S^\dagger 1,$$

with S the $[1 \times N_k N_l N_m]$-matrix whose components are given by $M(0,0,0)\,B_{s_k s_l s_m}(k,l,m)$. Akin to the previous derivations (46) this matrix-equation can be rewritten as

$$(\lambda R + \mu I)\,T_{new} = S^\dagger 1.$$

In particular our B-spline basis algorithm is a new algorithm that can be used for the resolvent (hypo-)elliptic diffusion process on $SE(2)$. The benefit over Fourier based algorithms is the local support of the basis functions, which allow for sparse representations.

In Fig. 6 we compare the impulse response for Tikhonov regularization via our B-spline expansion algorithm with the Brownian motion prior on $SE(2)$ (using a fine B-spline basis) to the exact solutions derived in [9], [10]. The strong accuracy of our algorithm shows that even in the discrete B-spline setting the probabilistic interpretation (Thm. 1) of our prior in $SE(2)$-template matching holds.

A.4 The Drunkman's Pencil

Similar to the diffusions on $\mathbb{R}^2$, given by (40), the hypo-elliptic diffusion process on $SE(2)$ is described by the following PDE:

$$\begin{cases} \frac{\partial}{\partial\tau}W &= (D_{\xi\xi}\partial_\xi^2 + D_{\theta\theta}\partial_\theta^2)W, \\ W|_{\tau=0} &= W_0, \end{cases} \quad (47)$$

initialized with $W_0 \in \mathbb{L}_2(\mathbb{R}^2)$ at time $\tau = 0$. The PDE can be used to obtain the solutions of our single patch problem by initializing W_0 with a smooth spike such as we did in Subsec. A.3, e.g. taking $W_0 = U_{f_1}(x,y,\theta) = B^n\left(\frac{x}{s_k}\right)B^n\left(\frac{y}{s_l}\right)B^n\left(\frac{\theta \bmod 2\pi}{s_m}\right)$.

The PDE in (47) is the forward Kolmogorov equation [65] of the following stochastic process [9]:

$$\begin{cases} \mathbf{x}(\tau) = \mathbf{x}(0) + \\ \qquad \sqrt{2D_{\xi\xi}}\,\epsilon_\xi\int_0^\tau(\cos\theta(\tau)\mathbf{e}_x + \sin\theta(\tau)\mathbf{e}_y)\frac{1}{2\sqrt{\tau}}\mathrm{d}\tau \\ \theta(\tau) = \theta(0) + \sqrt{\tau}\sqrt{2D_{\theta\theta}}\,\epsilon_\theta, \qquad \epsilon_\xi,\epsilon_\theta\,\mathcal{N}(0,1), \end{cases} \quad (48)$$

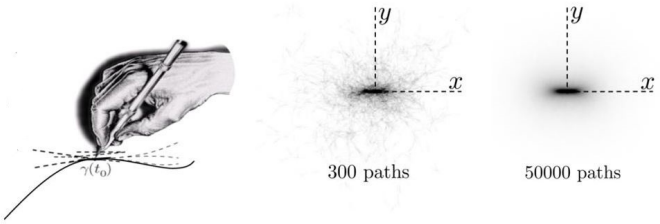

Fig. 7. Stochastic random process for contour enhancement.

where ϵ_ξ and ϵ_θ are sampled from a normal distribution with expectation 0 and unit standard deviation. The stochastic process in (48) can be interpreted as the motion of a drunkman's pencil: it randomly moves forward and backwards, and randomly changes its orientation along the way. The resolvent hypo-elliptic diffusion kernels $R_{\alpha,s}(g)$ (solutions to the fundamental single patch problem, up to scalar multiplication) can then also be obtained via Monte Carlo simulations, where the stochastic process is sampled many times with a negatively exponentially distributed traveling time ($P(\mathcal{T}=\tau)=\alpha e^{-\alpha\tau}$) in order to be able to estimate the probability density kernel $R_{\alpha,s}(g)$. This process is illustrated in Fig. 7.

APPENDIX B
THE SMOOTHING REGULARIZATION MATRIX R

When expanding the templates t and T in a finite B-Spline basis (Sec. 2 and 3 of the main article), the energy functionals (7), (11), (28) and (30) of the main article can be expressed in matrix vector form. The following theorems summarize how to compute the matrix R, which encodes the smoothing prior, for respectively the $\mathbb{R}^2$ and $SE(2)$ case.

Lemma 1. *The discrete smoothing regularization-term of energy functional (7) of the main article can be expressed directly in the B-Spline coefficients* $\mathbf{c}$ *as follows*

$$\iint_{\mathbb{R}^2} \|\nabla t(x,y)\|^2 \mathrm{d}x\mathrm{d}y = \mathbf{c}^\dagger R \mathbf{c}, \tag{49}$$

with $\mathbf{c}$ *given by Eq. (16) of the main article, and with*

$$R = R_x^k \otimes R_x^l + R_y^k \otimes R_y^l, \tag{50}$$

a $[N_k N_l \times N_k N_l]$ *matrix. The elements of the matrices in (50) are given by*

$$\begin{array}{ll} R_x^k(k,k') & = -\frac{1}{s_k}\frac{\partial^2 B^{2n+1}}{\partial x^2}(k'-k) \\ R_x^l(l,l') & = s_l B^{2n+1}(l'-l), \\ R_y^k(k,k') & = s_k B^{2n+1}(k'-k), \\ R_y^l(l,l') & = -\frac{1}{s_l}\frac{\partial^2 B^{2n+1}}{\partial y^2}(l'-l). \end{array} \tag{51}$$

Proof. For the sake of readability we divide the regularization-term in two parts:

$$\begin{aligned} \iint_{\mathbb{R}^2}\|\nabla t(x,y)\|^2\mathrm{d}x\mathrm{d}y &= \iint_{\mathbb{R}^2}\left|\tfrac{\partial t}{\partial x}(x,y)\right|^2 \\ &\quad + \left|\tfrac{\partial t}{\partial y}(x,y)\right|^2 \mathrm{d}x\mathrm{d}y, \end{aligned} \tag{52}$$

where

$$\begin{array}{ll} \mathcal{R}_x & = \iint_{\mathbb{R}^2}\left|\tfrac{\partial t}{\partial x}(x,y)\right|^2 \mathrm{d}x\mathrm{d}y, \text{ and} \\ \mathcal{R}_y & = \iint_{\mathbb{R}^2}\left|\tfrac{\partial t}{\partial y}(x,y)\right|^2 \mathrm{d}x\mathrm{d}y. \end{array}$$

We first derive the matrix-vector representation of $\mathcal{R}_x$ as follows:

$$\begin{aligned} \mathcal{R}_x &= \iint_{\mathbb{R}^2}\left|\tfrac{\partial t}{\partial x}(x,y)\right|^2 \mathrm{d}x\mathrm{d}y \\ &= \sum_{k,k'=1}^{N_k}\sum_{l,l'=1}^{N_l} \iint_{\mathbb{R}^2} \overline{c_{k,l}\tfrac{\partial B^n}{\partial x}(\tfrac{x}{s_k}-k)B^n(\tfrac{y}{s_l}-l)} \\ &\qquad c_{k,l}\tfrac{\partial B^n}{\partial x}(\tfrac{x}{s_k}-k')B^n(\tfrac{y}{s_l}-l')\mathrm{d}x\mathrm{d}y \\ &= \sum_{k,k'=1}^{N_k}\sum_{l,l'=1}^{N_l} \overline{c_{k,l}}c_{k',l'} \\ &\qquad \left[\int_{-\infty}^{\infty}\tfrac{\partial B^n}{\partial x}(\tfrac{x}{s_k}-k)\tfrac{\partial B^n}{\partial x}(\tfrac{x}{s_k}-k')\mathrm{d}x\right] \\ &\qquad \left[\int_{-\infty}^{\infty}B^n(\tfrac{y}{s_l}-l)B^n(\tfrac{y}{s_l}-l')\mathrm{d}y\right] \\ &\overset{1}{=} \sum_{k,k'=1}^{N_k}\sum_{l,l'=1}^{N_l} \overline{c_{k,l}}c_{k',l'}\left[\tfrac{1}{s_k}\left(\tfrac{\partial B^n}{\partial x}*\tfrac{\partial B^n}{\partial x}\right)(k'-k)\right] \\ &\qquad \left[s_l\left(B^n*B^n\right)(l'-l)\right] \\ &\overset{2}{=} \sum_{k,k'=1}^{M}\sum_{l,l'=1}^{N} \overline{c_{k,l}}c_{k',l'}\left[\tfrac{1}{s_k}\tfrac{\partial^2 B^{2n+1}}{\partial x^2}(k'-k)\right] \\ &\qquad \left[s_l B^{2n+1}(l'-l)\right] \\ &= \mathbf{c}^\dagger(R_x^k \otimes R_x^l)\mathbf{c}. \end{aligned} \tag{53}$$

Here the following properties are used:

1) The integrals of shifted B-splines can be expressed as convolutions:

$$\begin{aligned} &\int_{-\infty}^{\infty}\frac{\partial B^n}{\partial x}\left(\frac{x}{s_k}-k\right)\frac{\partial B^n}{\partial x}\left(\frac{x}{s_k}-k'\right)\mathrm{d}x \\ &\quad = -\frac{1}{s_k}\int_{-\infty}^{\infty}\frac{\partial B^n}{\partial u}(u)\frac{\partial B^n}{\partial u}((k'-k)-u)\mathrm{d}u \\ &\quad = -\frac{1}{s_k}\left(\frac{\partial B^n}{\partial u}*\frac{\partial B^n}{\partial u}\right)(k'-k). \end{aligned}$$

This is easily verified by substitution of integration variable ($u=-\frac{x}{s_k}+k$) and noting that $B^n(x)=B^n(-x)$ and $\frac{\partial B^n}{\partial x}(x)=-\frac{\partial B^n}{\partial x}(-x)$.

2) Convolution of two B-splines B^n of order n results in a B-Spline B^{2n+1} of order $2n+1$:

$$B^n * B^n = B^{2n+1}.$$

The elements of the matrices R_y^k and R_y^l are derived in a similar manner. □

As a result of Lemma 1 we can state the following.

Corollary 1. *Let* $V=\mathrm{span}\{B_{k,l}^n\}$*, with* $k=1,\ldots,N_k, l=1,\ldots,N_l$*, and shifted B-splines* $B_{k,l}^n$ *of order* n *(cf. Subsec. 2.4 of the main article). Let the energy function* $E_{lin}^B:\mathbb{R}^{N_kN_l}\to\mathbb{R}^+$ *be given by Eq. (13) of the main article. Then the optimal continuous template of the constrained optimization problem (cf. Eq. (7) of the main article)*

$$t^* = \underset{t\in V}{\mathrm{argmin}}\ E_{lin}(t)$$

has coefficients $\mathbf{c}^*$ *w.r.t. the B-spline basis for* V*, that are the unique solution of*

$$\nabla_{\mathbf{c}} E^B(\mathbf{c}^*) = \mathbf{0},$$

which boils down to Eq. (14) of the main article.

Lemma 2. *The discrete regularization-term of energy functional (28) of the main article can be expressed directly in the B-Spline coefficients:*

$$\iiint_{SE(2)} \|\nabla T\|_D \, \mathrm{d}x\mathrm{d}y\mathrm{d}\theta = \mathbf{c}^T (D_{\xi\xi} R_\xi + D_{\eta\eta} R_\eta + D_{\theta\theta} R_\theta)\mathbf{c}. \tag{54}$$

Matrix R_ξ is given by

$$\begin{aligned} R_\xi = &\left(R_\xi^{Ix} \otimes R_\xi^{Iy} \otimes R_\xi^{I\theta}\right) + \left(R_\xi^{IIx} \otimes R_\xi^{IIy} \otimes R_\xi^{II\theta}\right) \\ &+ \left(R_\xi^{IIIx} \otimes R_\xi^{IIIy} \otimes R_\xi^{III\theta}\right) + \left(R_\xi^{IVx} \otimes R_\xi^{IVy} \otimes R_\xi^{IV\theta}\right) \end{aligned} \tag{55}$$

with the elements of the matrices used in the Kronecker products given by

$$\begin{aligned} R_\xi^{Ix}(k,k') &= -\tfrac{1}{s_k} \tfrac{\partial^2 B^{2n+1}}{\partial x^2}(k'-k), \\ R_\xi^{Iy}(l,l') &= s_l B^{2n+1}(l'-l), \\ R_\xi^{I\theta}(m,m') &= \int_0^\pi \cos^2(\theta) B^n(\tfrac{\theta}{s_m} - m) B^n(\tfrac{\theta}{s_m} - m')\mathrm{d}\theta, \end{aligned} \tag{56}$$

$$\begin{aligned} R_\xi^{IIx}(k,k') &= -R_\xi^{IIIx}(k,k') = \tfrac{\partial B^{2n+1}}{\partial x}(k'-k), \\ R_\xi^{IIy}(l,l') &= -R_\xi^{IIIy}(l,l') = -\tfrac{\partial B^{2n+1}}{\partial y}(l'-l), \\ R_\xi^{II\theta}(m,m') &= R_\xi^{III\theta}(m,m') = \\ &\int_0^\pi \cos(\theta)\sin(\theta) B^n(\tfrac{\theta}{s_m} - m) B^n(\tfrac{\theta}{s_m} - m')\mathrm{d}\theta, \end{aligned} \tag{57}$$

$$\begin{aligned} R_\xi^{IVx}(k,k') &= s_k B^{2n+1}(k'-k), \\ R_\xi^{IVy}(l,l') &= -\tfrac{1}{s_l} \tfrac{\partial^2 B^{2n+1}}{\partial y^2}(l'-l), \\ R_\xi^{IV\theta}(m,m') &= \int_0^\pi \sin^2(\theta) B^n(\tfrac{\theta}{s_m} - m) B^n(\tfrac{\theta}{s_m} - m')\mathrm{d}\theta. \end{aligned} \tag{58}$$

Matrix R_η is given by

$$\begin{aligned} R_\eta = &\left(R_\xi^{IIx} \otimes R_\xi^{IIy} \otimes R_\xi^{IV\theta}\right) - \left(R_\xi^{IIx} \otimes R_\xi^{IIy} \otimes R_\xi^{II\theta}\right) \\ &- \left(R_\xi^{IIIx} \otimes R_\xi^{IIIy} \otimes R_\xi^{III\theta}\right) + \left(R_\xi^{IVx} \otimes R_\xi^{IVy} \otimes R_\xi^{I\theta}\right). \end{aligned} \tag{59}$$

Matrix R_θ is given by

$$R_\theta = \left(R_\theta^x \otimes R_\theta^y \otimes R_\theta^\theta\right), \tag{60}$$

with the elements of the matrices given by

$$\begin{aligned} R_\theta^x(k,k') &= s_k B^{2n+1}(k'-k), \\ R_\theta^y(l,l') &= s_l B^{2n+1}(l'-l), \\ R_\theta^\theta(m,m') &= -\tfrac{1}{s_m} \tfrac{\partial^2 B^{2n+1}}{\partial \theta^2}(m'-m). \end{aligned} \tag{61}$$

Proof. The proof of Lemma 2 follows the same steps as in the proof of Lemma 1, only here left-invariant derivatives are used (cf. Eq. (24) of the main article). The four separate terms $I - IV$ of Eq. (55) arise from the left invariant derivative ∂_ξ: $\left|\frac{\partial T}{\partial \xi}\right|^2 = \left|\cos(\theta)\frac{\partial T}{\partial x} + \sin(\theta)\frac{\partial T}{\partial y}\right|^2$. □

Lemma 2 has the following consequence.

Corollary 2. *Let $V = \mathrm{span}\{B^n_{k,l,m}\}$, with $k = 1, \ldots, N_k, l = 1, \ldots, N_l, m = 1, \ldots, N_m$, and shifted B-splines $B^n_{k,l,m}$ of order n (cf. Subsec. 3.5 of the main article). Let the energy function $\mathcal{E}^B_{lin} : \mathbb{R}^{N_k N_l N_m} \to \mathbb{R}^+$ be given by*

$$\mathcal{E}^B_{lin}(\mathbf{c}) = \frac{1}{N}\|S\mathbf{c} - \mathbf{y}\|^2 + \mathbf{c}^\dagger(\lambda R + \mu I)\mathbf{c}$$

With S and $\mathbf{y}$ given by (33) of the main article and with $R = D_{\xi\xi}R_\xi + D_{\eta\eta}R_\eta + D_{\theta\theta}R_\theta$ given in Lemma 2. Then the optimal continuous template of the constrained optimization problem (cf. Eq. (28) of the main article)

$$T^* = \operatorname*{argmin}_{T \in V} \ \mathcal{E}_{lin}(T)$$

has coefficients $\mathbf{c}^*$ w.r.t. the B-spline basis for V that are the unique solution of

$$\nabla_{\mathbf{c}} \mathcal{E}^B(\mathbf{c}^*) = \mathbf{0}.$$

This boils down to Eq. (14) of the main article, but then on $\mathbb{R}^{N_k N_l N_m}$ with matrices R and S given above.

APPENDIX C
NORMALIZED CROSS CORRELATION

In most applications it is necessary to make the detection system invariant to local contrast and luminosity changes. In our template matching framework this can either be achieved via certain pre-processing steps that filter out these variations, or by means of normalized cross-correlation. In normalized cross-correlation, both the template as well as the image are (locally) normalized to zero mean and unit standard deviation (with respect to the inner product used in the cross-correlations). In this section, we explain the necessary adaptations to extend the standard cross-correlation based framework to normalized cross-correlations.

C.1 Normalized Cross-Correlation in $\mathbb{R}^2$

In the usual cross-correlation based template matching approach, as described in Sec. 2 of the main article, we rely on the standard $\mathbb{L}_2(\mathbb{R}^2)$ inner product (Eq. (6) of the main article). In normalized cross-correlation it is however convenient to extend this inner product to include a windowing function m which indicates the relevant region (support) of the template. As such, the inner product with respect to windowing function m is given by

$$(t, f)_{\mathbb{L}_2(\mathbb{R}^2, m d\tilde{\mathbf{x}})} := \int_{\mathbb{R}^2} \overline{t(\tilde{\mathbf{x}})} f(\tilde{\mathbf{x}}) m(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}, \tag{62}$$

with associated norm $\|\cdot\|_{\mathbb{L}_2(\mathbb{R}^2, m d\tilde{\mathbf{x}})} = \sqrt{(\cdot,\cdot)_{\mathbb{L}_2(\mathbb{R}^2, m d\tilde{\mathbf{x}})}}$. The windowing function has to be a smooth function $m : \mathbb{R}^2 \to \mathbb{R}^+$ with $\int_{\mathbb{R}^2} m(\tilde{\mathbf{x}})\mathrm{d}\tilde{\mathbf{x}} = 1$. In this work, the use of a window m is also convenient to deal with boundary conditions in the optimization problems for template construction. We define

$$m(\mathbf{x}) := \varsigma \, e^{-\frac{\|\mathbf{x}\|^2}{s}} \sum_{i=0}^{n} \frac{(\|\mathbf{x}\|^2/s)^i}{i!}, \tag{63}$$

which smoothly approximates the indicator function $1_{[0,r]}(\|\mathbf{x}\|)$, covering a disk with radius r, when setting $s = \frac{2r^2}{1+2n}$, see e.g. [7, Fig. 2]. The constant ς normalizes the function such that $\int_{\mathbb{R}^2} m(\tilde{\mathbf{x}})\mathrm{d}\tilde{\mathbf{x}} = 1$.

In normalized cross-correlation the image is locally normalized (at position $\mathbf{x}$) to zero mean and unit standard deviation, which is done as follows

$$\hat{f}_{\mathbf{x}}(\tilde{\mathbf{x}}) := \frac{f(\tilde{\mathbf{x}}) - \langle f \rangle_{\mathcal{T}_{\mathbf{x}} m}}{\|f(\tilde{\mathbf{x}}) - \langle f \rangle_{\mathcal{T}_{\mathbf{x}} m}\|_{\mathbb{L}_2(\mathbb{R}^2, \mathcal{T}_{\mathbf{x}} m d\tilde{\mathbf{x}})}}, \tag{64}$$

with local mean $\langle f \rangle_m = (1, f)_{\mathbb{L}_2(\mathbb{R}^2, m\mathrm{d}\tilde{\mathbf{x}})}$. Template $\hat{t}$ can be obtained via normalization of a given template t via

$$\hat{t}(\tilde{\mathbf{x}}) := \frac{t(\tilde{\mathbf{x}}) - \langle t \rangle_m}{\|t(\tilde{\mathbf{x}}) - \langle t \rangle_m\|_{\mathbb{L}_2(\mathbb{R}^2, m\mathrm{d}\tilde{\mathbf{x}})}}. \tag{65}$$

Template matching is then done in the usual way (via (4) of the main article), however now $\hat{t}$ and $\hat{f}_{\mathbf{x}}$ are used instead of t and f. In fact, the entire $\mathbb{R}^2$ cross-correlation template matching, and template optimization framework is extended to normalized cross-correlation by substituting all instances of t with $\hat{t}$, f with $\hat{f}_{\mathbf{x}}$, and $(\cdot,\cdot)_{\mathbb{L}_2(\mathbb{R}^2)}$ with $(\cdot,\cdot)_{\mathbb{L}_2(\mathbb{R}^2, m d\tilde{\mathbf{x}})}$ in Sec. 2 of the main article. However, since templates $\hat{t}$ are directly constructed via the minimization of energy functionals, we will not explictely normalize the templates, unless they are obtained by other methods. E.g., Eq. (65) is used in the main article to construct basic templates obtained by averaging positive object patches (Subsec. 4.1 of the main article).

C.2 Normalized Cross-Correlation in $SE(2)$

Similar to the $\mathbb{R}^2$ case, templates and orientation scores are locally normalized to zero mean and unit standard deviation, however, now with respect to the $\mathbb{L}_2(SE(2), Md\tilde{g})$ inner product, which is given by

$$(T, U_f)_{\mathbb{L}_2(SE(2), M\mathrm{d}\tilde{g})} := \int_{\mathbb{R}^2} \int_0^{2\pi} \overline{T(\tilde{\mathbf{x}}, \tilde{\theta})} U_f(\tilde{\mathbf{x}}, \tilde{\theta}) M(\tilde{\mathbf{x}}, \tilde{\theta}) \mathrm{d}\tilde{\mathbf{x}} \mathrm{d}\tilde{\theta}, \tag{66}$$

with norm $\|\cdot\|_{\mathbb{L}_2(SE(2), M\mathrm{d}\tilde{g})} = \sqrt{(\cdot,\cdot)_{\mathbb{L}_2(SE(2), M\mathrm{d}\tilde{g})}}$. Also here windowing function M indicates the support of the template, and has the property $\int_{\mathbb{R}^2} \int_0^{2\pi} M(\tilde{\mathbf{x}}, \tilde{\theta}) \mathrm{d}\tilde{\mathbf{x}} \mathrm{d}\tilde{\theta} = 1$. We define

$$M(\mathbf{x}, \theta) := \frac{1}{2\pi} m(\mathbf{x}), \tag{67}$$

independent of θ and with $m(\mathbf{x})$ given by (63).

The (locally at g) normalized orientation score and template T are then given by

$$\hat{U}_{f,g}(\tilde{\mathbf{x}}, \tilde{\theta}) := \frac{U_f(\tilde{\mathbf{x}}, \tilde{\theta}) - \langle U_f \rangle_{\mathcal{L}_g M}}{\|U_f(\tilde{\mathbf{x}}, \tilde{\theta}) - \langle U_f \rangle_{\mathcal{L}_g M}\|_{\mathbb{L}_2(SE(2), \mathcal{L}_g M d\tilde{g})}}, \tag{68}$$

$$\hat{T}(\tilde{\mathbf{x}}, \tilde{\theta}) := \frac{T(\tilde{\mathbf{x}}, \tilde{\theta}) - \langle T \rangle_M}{\|T(\tilde{\mathbf{x}}, \tilde{\theta}) - \langle T \rangle_M\|_{\mathbb{L}_2(SE(2), M\mathrm{d}\tilde{g})}}, \tag{69}$$

with mean $\langle U_f \rangle_M = (1, U_f)_{\mathbb{L}_2(SE(2), M\mathrm{d}\tilde{g}))}$.

C.3 Efficient local normalization of $\hat{f}_{\mathbf{x}}$ and $\hat{U}_{f,g}$.

Since the normalized image $\hat{f}_{\mathbf{x}}$ depends on the location $\mathbf{x}$ it needs to be calculated for every translation of the template, which makes normalized cross-correlation computationally expensive. Therefore, (64) can be approximated by assuming that the local average is approximately constant in the area covered by m. That is, assuming $\langle f \rangle_{\mathcal{T}_{\mathbf{x}} m}(\tilde{\mathbf{x}}) \approx \langle f \rangle_{\mathcal{T}_{\tilde{\mathbf{x}}} m}(\tilde{\mathbf{x}}) = (m \star f)(\tilde{\mathbf{x}})$ for $\|\tilde{\mathbf{x}} - \mathbf{x}\| < r$, with r the radius that determines the extent of m, (64) is approximated as follows:

$$\hat{f}_{\mathbf{x}}(\tilde{\mathbf{x}}) \approx \frac{f(\tilde{\mathbf{x}}) - (m \star f)(\tilde{\mathbf{x}})}{\sqrt{(m \star (f - (m \star f))^2)(\tilde{\mathbf{x}})}}. \tag{70}$$

Similarly, in the $SE(2)$-case (68) can be approximated via

$$\hat{U}_{f,g}(\tilde{\mathbf{x}}, \tilde{\theta}) \approx \frac{U_f(\tilde{\mathbf{x}}, \tilde{\theta}) - (M \star_{SE(2)} U_f)(\mathbf{x}, \tilde{\theta})}{\sqrt{(M \star_{SE(2)} (U_f - (M \star_{SE(2)} U_f))^2)(\mathbf{x}, \tilde{\theta})}}. \tag{71}$$

C.4 Including a Region of Interest Mask

Depending on the application, large portions of the image might be masked out. This for example the case in retinal images (see circular masks in Fig. 9). To deal with this, template matching is only performed inside the region of interest defined by a mask image $m^{roi} : \mathbb{R}^2 \to \{0, 1\}$. Including such a mask is important in normalized template matching, and can be done by replacing the standard inner products by

$$(t, f)^{roi}_{\mathbb{L}_2(\mathbb{R}^2, m, d\tilde{\mathbf{x}})} := \frac{(t, f m^{roi})_{\mathbb{L}_2(\mathbb{R}^2, m, d\tilde{\mathbf{x}})}}{(1, m^{roi})_{\mathbb{L}_2(\mathbb{R}^2, m, d\tilde{\mathbf{x}})}}, \tag{72}$$

$$(T, U_f)^{roi}_{\mathbb{L}_2(SE(2), M, d\tilde{g})} := \frac{(T, U_f M^{roi})_{\mathbb{L}_2(SE(2), M, d\tilde{g})}}{(1, M^{roi})_{\mathbb{L}_2(SE(2), M, d\tilde{g})}}, \tag{73}$$

with $M^{roi}(\mathbf{x}, \theta) = m^{roi}(\mathbf{x})$.

APPENDIX D
ADDITIONAL DETAILS ON THE DETECTION PROBLEMS

In this section we describe additional details about the implementation and results of the three detection problems discussed in the main article.

D.1 Training Samples

In all three applications training samples were used to compute the templates. Positive training samples were centered around the object of interest. Negative training samples were centered around random locations in the image, but not within a certain distances to the true positive object location. In the retinal applications this distance was one optic disk radius, in the pupil detection application this was a normalized distance of 0.1 (cf. Eq.(39) of the main article). An selection of the 2D image pathes that were used in the experiments are shown in Fig. 8.

D.2 Processing Pipeline, Settings and Timings

D.2.1 Processing Pipeline

In all three application the same processing pipeline was used. The pipeline can be divided into the following 5 steps:

1) *Resizing*. Each input image is resized to a certain operating resolution and cropped to remove large regions with value 0 (outside the field of view mask in retinal images, see e.g. Fig. 9). The retinal images are resized such that the pixel size was approximately $40\mu m/pix$. In the pupil detection application no rescaling or cropping was done.
2) $\mathbb{R}^2$*-Processing*. In all three applications we applied a local intensity and contrast normalization step using an adaptation of [47] which we explain below. The

Fig. 8. A selection of positive and negative image patches f_i used in the training of templates.

TABLE 4
Average processing times. For optic nerve head detection (ONH) the average is taken over 1529 images of the TC, MESSIDOR, DRIVE and STARE database. For fovea detection the average is taken over 1408 images of the TC and MESSIDOR database. For pupil detection the average is taken over 1521 images of the BioID database.

	ONH		Fovea		Pupil (left & right)	
	$\mathbb{R}^2$	$SE(2)$	$\mathbb{R}^2$	$SE(2)$	$\mathbb{R}^2$	$SE(2)$
	Timings (ms)					
1. Rescaling	106	106	111	111	0	0
2. $\mathbb{R}^2$-Processing	66	66	64	64	71	71
3. OS Transform	0	108	0	108	0	121
4. $SE(2)$-Processing	0	5	0	5	0	6
5. Template Matching	20	195	19	190	26	116
Total	192	479	195	477	97	313
	Combined Total Timings (ms) - $\mathbb{R}^2$ and $SE(2)$					
	497		501		420	
	Combined Total Timings (ms) - Fovea and ONH					
	730					

TABLE 5
Success rates for optic nerve head detection ($\pm$ standard deviation, number of fails in parenthesis) with varying accuracy requirements in 5-fold cross validation. Maximum distance to ground truth location is expressed in optic disk radius R.

Database (# of images)		Maximum distance to ground truth R/8	R/4	R/2	R	2R
ES (SLO)	(208)	98.05% $\pm$ 2.04% (4)	100.0% $\pm$ 0.00% (0)	100.0% $\pm$ 0.00% (0)	100.0% $\pm$ 0.00% (0)	100.0% $\pm$ 0.00% (0)
TC	(208)	84.19% $\pm$ 4.34% (33)	94.54% $\pm$ 3.51% (11)	99.52% $\pm$ 1.06% (1)	100.0% $\pm$ 0.00% (0)	100.0% $\pm$ 0.00% (0)
MESSIDOR	(1200)	73.07% $\pm$ 3.69% (323)	94.41% $\pm$ 1.47% (67)	99.50% $\pm$ 0.46% (6)	99.92% $\pm$ 0.19% (1)	100.0% $\pm$ 0.00% (0)
DRIVE	(40)	70.84% $\pm$ 26.0% (13)	91.69% $\pm$ 12.3% (4)	98.18% $\pm$ 4.07% (1)	98.18% $\pm$ 4.07% (1)	100.0% $\pm$ 0.00% (0)
STARE	(81)	48.12% $\pm$ 10.27% (42)	74.94% $\pm$ 6.52% (20)	89.39% $\pm$ 8.16% (9)	98.67% $\pm$ 2.98% (1)	98.67% $\pm$ 2.98% (1)
All Images	(1737)	76.11% $\pm$ 2.58% (415)	94.13% $\pm$ 0.79% (102)	99.02% $\pm$ 0.26% (17)	99.83% $\pm$ 0.26% (3)	99.94% $\pm$ 0.13% (1)

locally normalized image $\hat{f}$ is then mapped through an error function via $\mathrm{erf}(8\hat{f})$ to dampen outliers.

3) *Orientation score transform*. The processed image is then taken as input for an orientation score transform using Eq. (23) of the main article. For the oriented wavelets we used cake wavelets [6], [7] of size $[51\times51]$ and with angular resolution $s_\theta = \pi/12$, and with sampling θ from 0 to π.
4) $SE(2)$*-Processing*. For phase-invariant, nonlinear, left-invariant [10], and contractive [48] processing on SE(2), we work with the modulus of the complex valued orientation scores rather than with the complex-valued scores themselves (taking the modulus of quadrature filter responses is an effective technique for line detection, see e.g. Freeman et al. [49]).
5) *Template Matching*. Finally we perform template matching using respectively Eqs. (3),(4) and (5) of the main article for the $\mathbb{R}^2$ case and Eqs. (3),(25) and (26) of the main article for the $SE(2)$ case.

Regarding the image resolutions (step 1) we note that the average image size after rescaling was $[300 \times 300]$. The average image resolutions in each database were as follows:

- *ES (SLO)* contained images of average resolution $13.9\mu m/pix$.
- *TC* contained images of average resolution $9.4\mu m/pix$.
- *MESSIDOR* contained images of 3 cameras with average resolutions $13.6\mu m/pix$, $9.1\mu m/pix$ and $8.6\mu m/pix$.
- *DRIVE* contained images of average resolution $21.9\mu m/pix$.
- *STARE* contained images of average resolution $17.6\mu m/pix$.

Regarding local image normalization (step 2) we note the following. Local image normalization was done using an adaptation of [47]. The method first computes a local average and standard deviation of pixel intensities, and the image is locally normalized to zero mean and unit standard deviation. This is done via Eq. (70). Then a background mask is construct by setting pixels with a larger distance than 1 standard deviation to the average (Mahalanobis distance) to 0, and other pixels to 1. This mask is then used to ignore outliers in a second computation of the local average and standard deviation. The final normalized image is again computed via Eq. (70) but now with the inclusion of the background mask, see Eq. (72).

D.2.2 Template Settings

In the retinal applications we used $\mathbb{R}^2$ templates of size $[N_x \times N_y] = [251 \times 251]$ which were covered by a grid of B-spline basis functions of size $[N_k \times N_l] = [51\times51]$, the $SE(2)$ templates were of size $[N_x \times N_y \times N_\theta] = [251 \times 251 \times 12]$ and were covered by a grid of B-spline basis functions of size $[N_k \times N_l \times N_m] = [51 \times 51 \times 12]$.

In the pupil detection application we used $\mathbb{R}^2$ templates of size $[N_x \times N_y] = [101 \times 101]$ which were also covered by a grid of B-spline basis functions of size $[N_k \times N_l] = [51 \times 51]$, the $SE(2)$ templates were of size $[N_x \times N_y \times N_\theta] = [101 \times 101 \times 12]$ and were also covered by a grid of B-spline basis functions of size $[N_k \times N_l \times N_m] = [51 \times 51 \times 12]$.

D.2.3 Timings

We computed the average time for detecting one (or two) object(s) in an image and tabulated the results in Tab. 4. Here we sub-divided the timings into the 5 processing steps explained in Subsec. D.2.1. The average (full) processing time on the retinal images was in both applications approximately $500ms$. When both the ONH and fovea are detected by the same processing pipeline the processing took $730ms$. For pupil detection the average time to detect *both* the left and right pupil on the *full* images was $420ms$.

The retinal images were on average of size $[1230\times1792]$, and $[300 \times 300]$ after cropping and resizing. The images in the pupil detection application were not resized or cropped and were of size $[286 \times 384]$.

All experiments were performed using Wolfram *Mathematica* 10.4, on a computer with an Intel Core i703612QM CPU and 8GB memory.

D.3 Detection Results

In this section we provide the results for the three separate applications. A general discussion of these results can be found in the main article.

D.3.1 Optic Nerve Head Detection

A Table of detection performance for each type of template is provided in Tab. 1 of the main article. In Fig. 9 we show the 3 failed cases for ONH detection, and a selection of correct ONH localizations in difficult images. In Table 5 we show detection results for varying accuracy criteria. Note

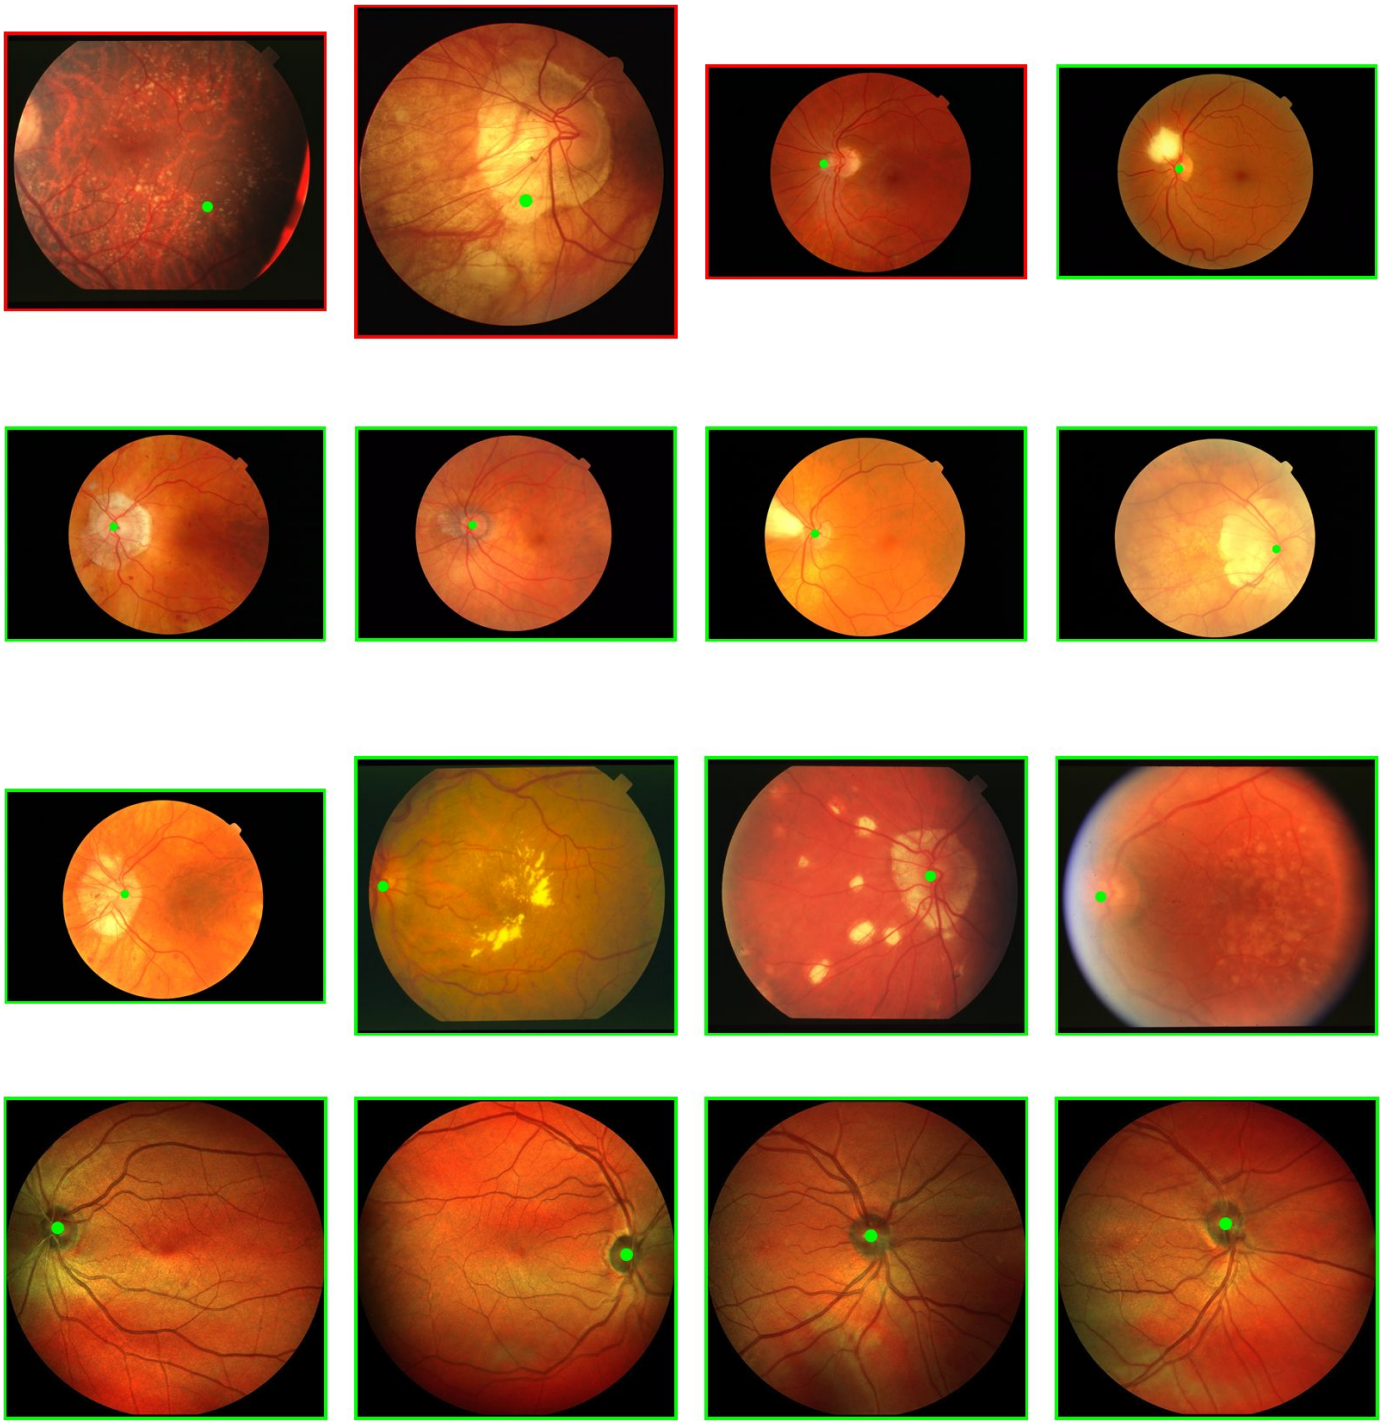

Fig. 9. Detection results of our best method for optic nerve head detection in retinal images. Successful detection are indicated with a green frame around the image, failed detections are indicated with a red frame. In the ONH detection application there were only 3 fails in a set of 1737 images.

that detection results are typically reported for the accuracy requirement of 1 optic disk radius with the target (see also state-of-the-art comparison in Table 2 of the main article).

D.3.2 Fovea Detection

A Table of detection performance for each type of template is provided in Tab. 6. In Fig. 10 we show next to a selection of successful detections the only 5 failed cases on images from conventional fundus (CF) cameras (TC, MESSIDOR, DRIVE, STARE), and 3 of the failed detections in images coming from an scanning laser ophthalmoscopy (SLO) camera.

As can also be read from Tab. 6, we found that fovea detection in SLO images was significantly more difficult than fovea detection in CF images. The reason for this is that on SLO images the clear dark blob-like shape is not always present on these images. Compare for example the positive fovea patches from Fig. 8 (where one generally sees a dark blob at the center) with the fovea locations in the bottom row of images in Figs. 9 and 10.

Additionally, the ES (SLO) and CF databases are also more difficult than the MESSIDOR database for fovea detection, as these two databases contain a mix of both fovea centered and ONH centered images. The MESSIDOR database contains only fovea centered images, in which case the fovea is always located around the center of the image. Therefore, even though MESSIDOR is one of the most used databases, it might not be the most representative database for fovea detection benchmarking.

We show detection performance for a range of accuracy requirements in Table 7 for the different databases used in our experiments, and in Table 8 a comparison to the state of the art. There we see that for the stricter requirement of detection within half an optic disk radius our method still outperforms the state of the art. We also see that with further decreasing the acceptance distance ($R/4$ or lower) none of the methods provided acceptable results.

D.3.3 Pupil Detection

A Table of detection performance for each type of template is provided in Tab. 9. In Fig. 11 we show a selection of failed and successful detections. By inspection of the failed cases we found that a main source of failed detections was due to rotations of the head. As stated in the previous section D.2 we did not employ a rotation invariant detection scheme. Doing so might improve the results. Other failed detections could be attributed to closed eyes, reflection of glasses, distracting background objects and different scales (object distance to camera).

TABLE 9
Average template matching results ($\pm$ standard deviation) for pupil detection in 5-fold cross validation, number of failed detections in parentheses. A successful detection has a normalized error $e \leq 0.1$.

Template ID	BioID (Full image) 1521	BioID (Periocular image) 1521
$\mathbb{R}^2$ templates		
$A_{\mathbb{R}^2}$	41.03% ± 1.45% (897)	59.70% ± 1.52% (613)
$B_{lin:\mathbb{R}^2}$	0.00% ± 0.00% (1521)	3.62% ± 1.09% (1466)
$C_{lin:\mathbb{R}^2}$	12.95% ± 2.22% (1324)	67.26% ± 2.55% (498)
$D_{lin:\mathbb{R}^2}$	8.28% ± 1.80% (1395)	75.68% ± 2.33% (370)
$E_{lin:\mathbb{R}^2}$	11.51% ± 2.25% (1346)	71.47% ± 2.76% (434)
$B_{log:\mathbb{R}^2}$	0.00% ± 0.00% (1521)	0.00% ± 0.00% (1521)
$C_{log:\mathbb{R}^2}$	12.89% ± 2.06% (1325)	39.91% ± 3.37% (914)
$D_{log:\mathbb{R}^2}$	1.84% ± 0.95% (1493)	22.09% ± 2.37% (1185)
$E_{log:\mathbb{R}^2}$	10.39% ± 2.26% (1363)	37.21% ± 4.37% (955)
$SE(2)$ templates		
$A_{SE(2)}$	57.72% ± 1.68% (643)	75.34% ± 1.31% (375)
$B_{lin:SE(2)}$	8.74% ± 2.00% (1388)	41.81% ± 5.04% (885)
$C_{lin:SE(2)}$	84.61% ± 4.19% (234)	86.78% ± 3.68% (201)
$D_{lin:SE(2)}$	85.53% ± 3.44% (220)	87.18% ± 3.71% (195)
$E_{lin:SE(2)}$	85.47% ± 3.82% (221)	87.11% ± 3.87% (196)
$B_{log:SE(2)}$	0.00% ± 0.00% (1521)	0.13% ± 0.29% (1519)
$C_{log:SE(2)}$	86.52% ± 0.77% (205)	93.95% ± 1.33% (92)
$D_{log:SE(2)}$	75.21% ± 2.18% (377)	89.48% ± 2.27% (160)
$E_{log:SE(2)}$	83.30% ± 1.68% (254)	92.77% ± 1.02% (110)
Template combinations (sorted on performance full image)		
*$C_{lin:\mathbb{R}^2} + E_{lin:SE(2)}$	93.49% ± 1.49% (99)	95.60% ± 1.46% (67)
$C_{lin:\mathbb{R}^2} + D_{lin:SE(2)}$	93.16% ± 1.54% (104)	95.00% ± 1.15% (76)
$E_{lin:\mathbb{R}^2} + E_{lin:SE(2)}$	93.10% ± 1.04% (105)	95.59% ± 0.89% (67)
$E_{lin:\mathbb{R}^2} + D_{lin:SE(2)}$	92.97% ± 1.62% (107)	95.27% ± 1.31% (72)
$C_{lin:\mathbb{R}^2} + C_{lin:SE(2)}$	92.70% ± 1.41% (111)	95.33% ± 0.97% (71)
$E_{lin:\mathbb{R}^2} + C_{lin:SE(2)}$	92.64% ± 0.94% (112)	95.33% ± 0.94% (71)
$D_{lin:\mathbb{R}^2} + D_{lin:SE(2)}$	92.51% ± 0.96% (114)	95.79% ± 0.82% (64)
$D_{lin:\mathbb{R}^2} + E_{lin:SE(2)}$	92.24% ± 1.23% (118)	95.86% ± 0.89% (63)
$E_{log:\mathbb{R}^2} + D_{lin:SE(2)}$	92.11% ± 2.26% (120)	93.23% ± 1.93% (103)
$D_{lin:\mathbb{R}^2} + C_{log:SE(2)}$	92.05% ± 1.52% (121)	95.14% ± 0.78% (74)
...		...
Template combinations (sorted on performance periocular image)		
*$D_{lin:\mathbb{R}^2} + E_{lin:SE(2)}$	92.24% ± 1.23% (118)	95.86% ± 0.89% (63)
$D_{lin:\mathbb{R}^2} + D_{lin:SE(2)}$	92.51% ± 0.96% (114)	95.79% ± 0.82% (64)
$D_{lin:\mathbb{R}^2} + C_{lin:SE(2)}$	91.52% ± 1.25% (129)	95.73% ± 0.77% (65)
$E_{lin:\mathbb{R}^2} + E_{lin:SE(2)}$	93.10% ± 1.04% (105)	95.59% ± 0.89% (67)
$C_{lin:\mathbb{R}^2} + E_{lin:SE(2)}$	93.49% ± 1.49% (99)	95.60% ± 1.46% (67)
$E_{lin:\mathbb{R}^2} + C_{lin:SE(2)}$	92.64% ± 0.94% (112)	95.33% ± 0.94% (71)
$C_{lin:\mathbb{R}^2} + C_{lin:SE(2)}$	92.70% ± 1.41% (111)	95.33% ± 0.97% (71)
$E_{lin:\mathbb{R}^2} + D_{lin:SE(2)}$	92.97% ± 1.62% (107)	95.27% ± 1.31% (72)
$D_{lin:\mathbb{R}^2} + E_{log:SE(2)}$	91.72% ± 1.23% (126)	95.27% ± 0.79% (72)
$D_{lin:\mathbb{R}^2} + C_{log:SE(2)}$	92.05% ± 1.52% (121)	95.14% ± 0.78% (74)
...		...
†$A_{\mathbb{R}^2} + A_{SE(2)}$	61.34% ± 1.54% (588)	68.18% ± 1.25% (484)
...		...

Best template combination that does not rely on logistic regression.
†*Best template combination that does not rely on template optimization.*

APPENDIX E
ROTATION-SCALE INVARIANT MATCHING

E.1 A Basic Extension

The extension to rotation and scale invariant object localization of the 2D cross-correlation based template matching approach, described in Eqs. (3)-(5) of the main article, is as follows. For the linear potential function (Eq. (4) of the main article) we can define

$$P^{\mathbb{R}^2}_{lin,inv}(\mathbf{x}) := \max_{\substack{a \in [a_-, a_+],\\ \alpha \in [0, 2\pi)}} (\mathcal{T}_{\mathbf{x}} \mathcal{S}_a \mathcal{R}_\alpha \, t, f)_{\mathbb{L}_2(\mathbb{R}^2)}, \quad (74)$$

TABLE 6
Average template matching results ($\pm$ standard deviation) for fovea detection in 5-fold cross validation, number of failed detections in parentheses.

Template ID	ES (SLO) 208	TC 208	MESSIDOR 1200	All Images 1616
		$\mathbb{R}^2$ templates		
$A_{\mathbb{R}^2}$	76.36% ± 6.79% (49)	98.24% ± 2.74% (3)	98.41% ± 0.22% (19)	95.60% ± 0.98% (71)
$B_{lin:\mathbb{R}^2}$	23.50% ± 3.81% (159)	31.66% ± 9.03% (142)	51.19% ± 5.97% (587)	45.07% ± 3.33% (888)
$C_{lin:\mathbb{R}^2}$	45.65% ± 8.61% (113)	98.24% ± 2.74% (3)	98.59% ± 0.36% (17)	91.77% ± 1.26% (133)
$D_{lin:\mathbb{R}^2}$	44.21% ± 4.62% (116)	99.49% ± 1.14% (1)	98.84% ± 0.31% (14)	91.90% ± 0.59% (131)
$E_{lin:\mathbb{R}^2}$	46.10% ± 8.11% (112)	98.86% ± 1.57% (2)	98.67% ± 0.34% (16)	91.95% ± 1.18% (130)
$B_{log:\mathbb{R}^2}$	1.43% ± 1.31% (205)	10.27% ± 5.09% (185)	20.07% ± 3.00% (959)	16.53% ± 2.52% (1349)
$C_{log:\mathbb{R}^2}$	9.59% ± 3.74% (188)	70.30% ± 8.57% (61)	77.61% ± 4.64% (267)	68.06% ± 3.53% (516)
$D_{log:\mathbb{R}^2}$	11.48% ± 4.70% (184)	83.47% ± 7.80% (32)	88.22% ± 2.81% (141)	77.90% ± 2.00% (357)
$E_{log:\mathbb{R}^2}$	2.86% ± 2.62% (202)	79.68% ± 7.92% (40)	84.79% ± 5.16% (181)	73.82% ± 2.62% (423)
		$SE(2)$ templates		
$A_{SE(2)}$	67.81% ± 4.69% (67)	79.13% ± 9.11% (40)	98.25% ± 0.68% (21)	92.08% ± 0.84% (128)
$B_{lin:SE(2)}$	83.19% ± 2.76% (35)	71.53% ± 7.36% (58)	91.31% ± 0.68% (104)	87.81% ± 1.25% (197)
$C_{lin:SE(2)}$	83.65% ± 3.18% (34)	84.13% ± 6.25% (32)	98.23% ± 1.04% (21)	94.62% ± 0.36% (87)
$D_{lin:SE(2)}$	73.57% ± 4.71% (55)	83.69% ± 6.83% (33)	97.88% ± 1.17% (25)	93.01% ± 1.09% (113)
$E_{lin:SE(2)}$	77.83% ± 4.29% (46)	84.88% ± 6.69% (30)	98.22% ± 1.23% (21)	94.00% ± 0.93% (97)
$B_{log:SE(2)}$	75.49% ± 5.73% (51)	60.80% ± 5.68% (80)	92.79% ± 1.98% (86)	86.56% ± 2.20% (217)
$C_{log:SE(2)}$	79.33% ± 6.57% (43)	70.87% ± 10.28% (59)	96.90% ± 0.71% (37)	91.39% ± 1.36% (139)
$D_{log:SE(2)}$	62.09% ± 6.66% (79)	72.57% ± 8.59% (54)	96.64% ± 1.05% (40)	89.30% ± 0.63% (173)
$E_{log:SE(2)}$	68.34% ± 8.59% (66)	72.20% ± 8.53% (55)	96.57% ± 0.96% (41)	89.98% ± 1.25% (162)
		Template combinations (sorted on performance)		
$C_{lin:\mathbb{R}^2} + C_{log:SE(2)}$	97.17% ± 3.01% (6)	99.17% ± 1.13% (2)	99.74% ± 0.38% (3)	99.32% ± 0.26% (11)
* $A_{\mathbb{R}^2} + C_{lin:SE(2)}$	98.08% ± 2.03% (4)	98.07% ± 1.95% (4)	99.68% ± 0.33% (4)	99.26% ± 0.47% (12)
$E_{lin:\mathbb{R}^2} + C_{log:SE(2)}$	96.20% ± 3.15% (8)	99.17% ± 1.13% (2)	99.75% ± 0.23% (3)	99.20% ± 0.35% (13)
$E_{lin:\mathbb{R}^2} + C_{lin:SE(2)}$	96.65% ± 2.13% (7)	99.17% ± 1.13% (2)	99.66% ± 0.36% (4)	99.19% ± 0.42% (13)
$C_{lin:\mathbb{R}^2} + C_{lin:SE(2)}$	97.14% ± 1.97% (6)	98.78% ± 1.78% (3)	99.58% ± 0.31% (5)	99.13% ± 0.40% (14)
$A_{\mathbb{R}^2} + E_{lin:SE(2)}$	97.59% ± 1.73% (5)	98.07% ± 1.95% (4)	99.59% ± 0.28% (5)	99.13% ± 0.25% (14)
$E_{lin:\mathbb{R}^2} + E_{lin:SE(2)}$	96.16% ± 2.76% (8)	99.17% ± 1.13% (2)	99.58% ± 0.31% (5)	99.07% ± 0.38% (15)
$E_{lin:\mathbb{R}^2} + D_{lin:SE(2)}$	95.71% ± 3.07% (9)	99.17% ± 1.13% (2)	99.58% ± 0.31% (5)	99.01% ± 0.40% (16)
$C_{lin:\mathbb{R}^2} + E_{lin:SE(2)}$	96.16% ± 2.76% (8)	98.78% ± 1.78% (3)	99.58% ± 0.31% (5)	99.01% ± 0.51% (16)
$A_{\mathbb{R}^2} + C_{log:SE(2)}$	96.65% ± 2.13% (7)	98.07% ± 1.95% (4)	99.58% ± 0.42% (5)	99.01% ± 0.26% (16)
...		...		...
† $A_{\mathbb{R}^2} + A_{SE(2)}$	92.85 % ± 4.68% (15)	95.84% ± 2.58% (8)	99.58% ± 0.30% (5)	98.27% ± 0.70% (28)
...		...		...

Best template combination that does not rely on logistic regression.
†*Best template combination that does not rely on template optimization.*

TABLE 7
Success rates for fovea detection ($\pm$ standard deviation, number of fails in parenthesis) with varying accuracy requirements in 5-fold cross validation. Maximum distance to ground truth location is expressed in optic disk radius R.

Database (# of images)		Maximum distance to ground truth R/8	R/4	R/2	R	2R
ES (SLO)	(208)	66.91% ± 4.64% (69)	92.85% ± 3.16% (15)	94.74% ± 1.93% (11)	97.17% ± 3.01% (6)	97.66% ± 3.28% (5)
TC	(208)	49.51% ± 4.07% (106)	80.33% ± 3.22% (40)	95.41% ± 1.77% (9)	99.17% ± 1.13% (2)	99.61% ± 0.88% (1)
MESSIDOR	(1200)	61.81% ± 2.64% (459)	90.56% ± 1.31% (113)	98.07% ± 0.87% (23)	99.74% ± 0.38% (3)	100.0% ± 0.00% (0)
All Images	(1616)	60.78% ± 1.84% (634)	89.60% ± 0.80% (168)	97.34% ± 0.65% (43)	99.32% ± 0.26% (11)	99.63% ± 0.40% (6)

TABLE 8
Success rates for fovea detection (number of fails in parenthesis) with varying accuracy requirements; a comparison to literature using the MESSIDOR database. Maximum distance to ground truth location is expressed in optic disk radius R.

Method	Maximum distance to ground truth				
	R/8	R/4	R/2	R	2R
Niemeijer et al. [28], [55]	75.67% (292)	93.50% (78)	96.83% (38)	97.92% (25)	-
Yu et al. et al. [54]	-	-	95.00% (60)	-	-
Gegundez-Arias et al. [28]	80.42% (235)	93.90% (73)	96.08% (47)	96.92% (37)	97.83% (26)
Giachetti et al. [45]	-	-	-	99.10% (11)	-
Aquino [53]	-	-	-	98.20% (21)	-
Proposed	61.81% (459)	90.56% (113)	98.07% (23)	99.74% (3)	100.0% (0)

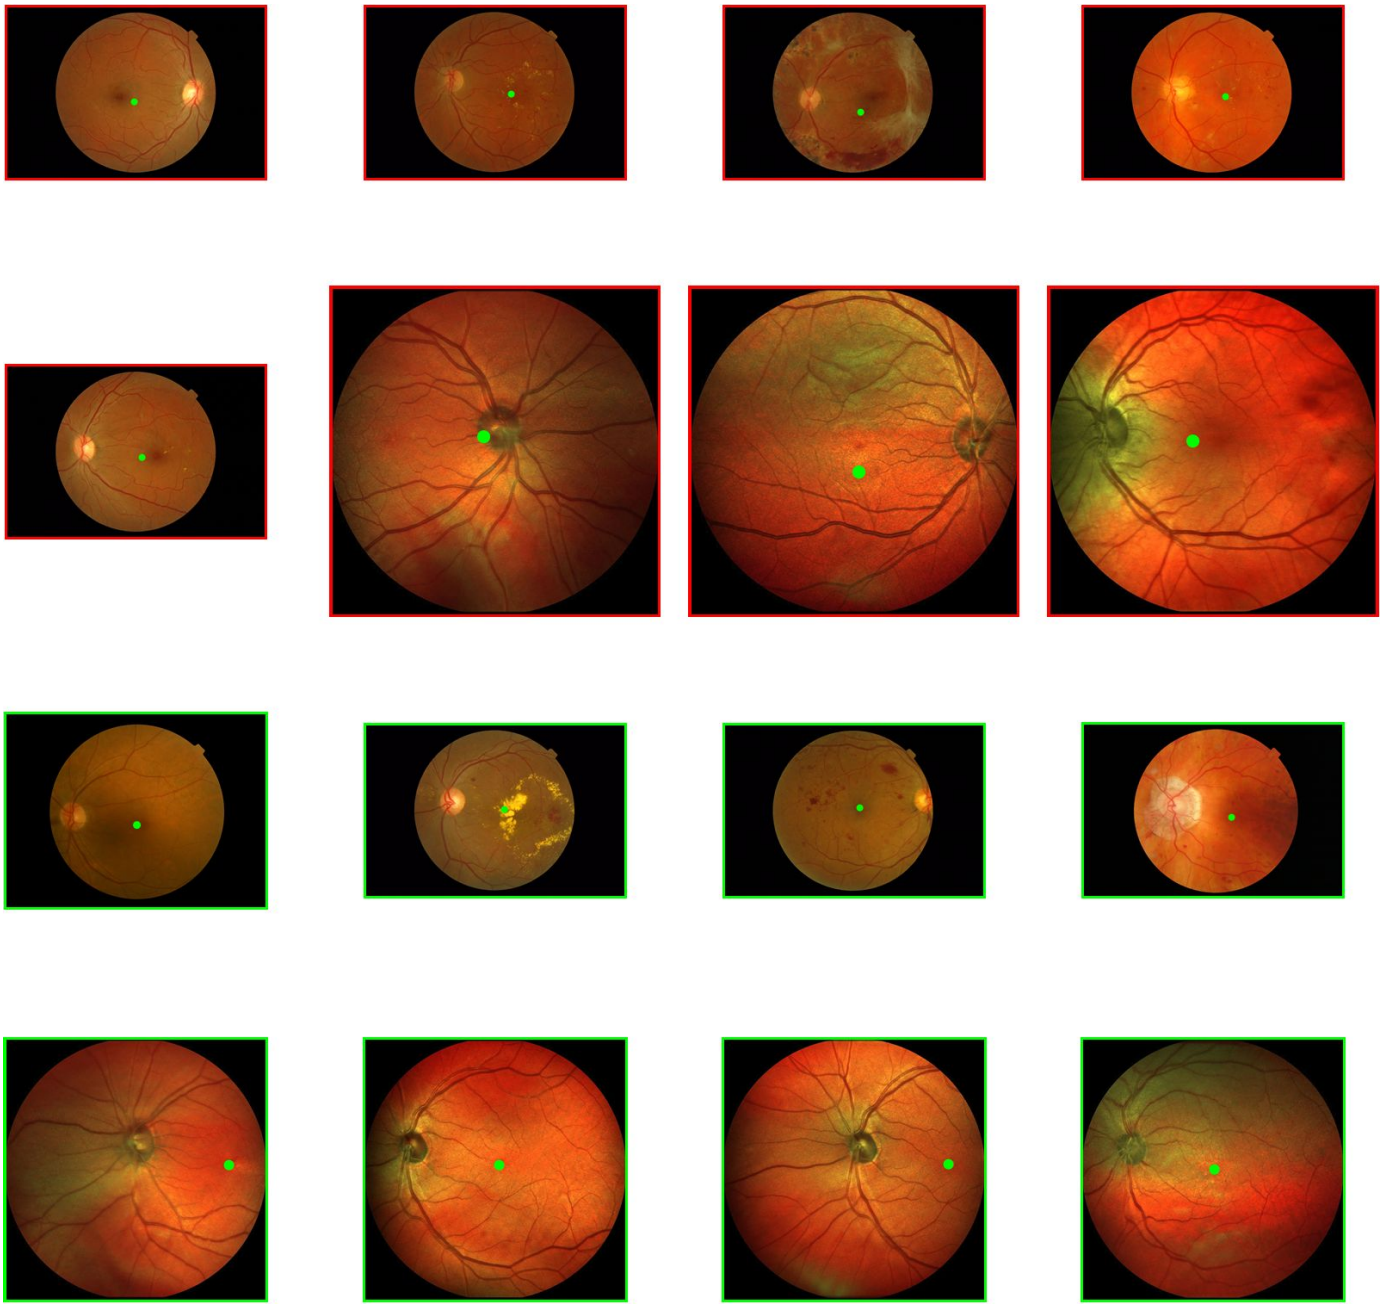

Fig. 10. Detection results of our best method for fovea detection in retinal images. Successful detection are indicated with a green frame around the image, failed detections are indicated with a red frame. In the fovea detection application there were only 5 fails in a set of 1408 conventional fundus (CF) camera images. Out of the 208 scanning laser ophthalmoscopy (SLO) images there were 6 fails, 3 of them are shown in this figure.

Fig. 11. Detection results of our best method for pupil detection. Successful detection are indicated with a green frame around the image, failed detections are indicated with a red frame.

and for the logistic regression case (Eq. (5) of the main article) we define

$$P^{\mathbb{R}^2}_{log,inv}(\mathbf{x}) := \max_{\substack{a \in [a_-, a_+],\\ \alpha \in [0, 2\pi)}} \sigma\left((\mathcal{T}_{\mathbf{x}}\mathcal{S}_a\mathcal{R}_\alpha\, t, f)_{\mathbb{L}_2(\mathbb{R}^2)}\right), \tag{75}$$

with σ the logistic sigmoid function defined in Eq. (5) of the main article, and with rotation operator $\mathcal{R}_\alpha$ and scaling operator $\mathcal{S}_a$ defined by

$$(\mathcal{R}_\alpha t)(\mathbf{x}) = t(\mathbf{R}_\alpha^{-1}\mathbf{x}), \tag{76}$$

$$(\mathcal{S}_a t)(\mathbf{x}) = a^{-1} t(a\mathbf{x}), \tag{77}$$

with rotation matrix $\mathbf{R}_\alpha$ representing a counter clockwise rotation of angle α. By taking the maximum over scales a (in a suitable range $[a_-, a_+]$) and rotations α, the response of the best matching template is obtained at each location $\mathbf{x}$, and invariance is obtained with respect to scaling and rotation of the object of interest.

The rotation/scale invariant extension of the $SE(2)$ cross-correlation template matching case (Eqs. (25)-(26) of the main article) is equally straightforward; for the linear potential we define

$$P^{SE(2)}_{lin,inv}(\mathbf{x}) := \max_{\substack{a \in [a_-, a_+],\\ \alpha \in [0, 2\pi)}} (\mathcal{T}_{\mathbf{x}}\mathcal{S}_a\mathcal{R}_\alpha\, T, U_f)_{\mathbb{L}_2(SE(2))}, \tag{78}$$

and for the logistic potential we define

$$P^{SE(2)}_{log,inv}(\mathbf{x}) := \max_{\substack{a \in [a_-, a_+],\\ \alpha \in [0, 2\pi)}} \sigma\left((\mathcal{T}_{\mathbf{x}}\mathcal{S}_a\mathcal{R}_\alpha\, T, U_f)_{\mathbb{L}_2(SE(2))}\right), \tag{79}$$

with for orientation score objects $T, U_f \in \mathbb{L}_2(SE(2))$ the rotation and scaling operators defined respectively by

$$(\mathcal{R}_\alpha T)(\mathbf{x}, \theta) = T(\mathbf{R}_\alpha^{-1}\mathbf{x}, \theta - \alpha), \tag{80}$$

$$(\mathcal{S}_a T)(\mathbf{x}, \theta) = a^{-1} T(a\mathbf{x}, \theta). \tag{81}$$

It depends on the addressed template matching problem whether or not such invariance is desirable or not. In many applications the object is to be found in a human environment context, in which some objects tend to appear in specific orientations or at typical scales, and in which case rotation/scale invariance might not be desirable. E.g. the sizes

of anatomical structures in the retina are relatively constant among different subjects (constant scale) and retinal images are typically taken at a fixed orientation (constant rotation). In the pupil detection problem the subjects typically appear in upright position behind the camera (constant rotation), and within a reasonable distance to the camera (constant scale). In the next Subsec. E.2.1 we indeed show that in the applications considered in this manuscript rotation/scale invariance is not necessarily a desired property, and that computation time linearly increases with the number of rotations/scalings tested for (cf. Subsec. E.2.2).

E.2 Results with Rotation and Scale Invariance

Here we perform rotation and scale invariant template matching via the extension described in Subsec. E.1. We selected the best template combination for each specific application and compared non-invariant template matching (as described in the main article) to rotation and/or scale invariant template matching (Subsec. E.1). The best template combination for ONH detection was $A_{\mathbb{R}^2} + C_{log:SE(2)}$, for fovea detection this was $C_{lin:\mathbb{R}^2} + C_{log:SE(2)}$, and for pupil detection this was $D_{lin:\mathbb{R}^2} + E_{lin:SE(2)}$.

For the retinal applications we only tested for rotation invariance with

$$\alpha \in \{-\frac{\pi}{6}, -\frac{\pi}{8}, -\frac{\pi}{12}, -\frac{\pi}{24}, 0, \frac{\pi}{24}, \frac{\pi}{12}, \frac{\pi}{8}, \frac{\pi}{6}\},$$

and did not included scale invariance since each retinal image was already rescaled to a standardized resolution (see Subsec. D.2.1). In pupil detection we tested for a range of scalings with

$$a \in \{0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3\}$$

to deal with varying pupil sizes caused by varying distances to the camera; and we tested for a range of rotations with

$$\alpha \in \{-\frac{\pi}{4}, -\frac{\pi}{8}, -\frac{\pi}{16}, 0, \frac{\pi}{16}, \frac{\pi}{8}, \frac{\pi}{4}\}$$

to deal with rotations of the head.

E.2.1 Detection Results

The detection results are shown in Table. 10. Here we can see that in all three applications the inclusion of a rotation/scale invariant matching scheme results in a slight decrease in performance. This can be explained by the fact that variations in scale an rotation within the databases are small, and that the trained templates can already deal robustly with these variations (due to the presence of such variations in the training set). By introducing rotation/scale invariance one then only increases the likelihood of false positive detections.

E.2.2 Computation Time

The effect on computation time of rotation/scale invariant matching is shown in Fig. 12. Here one sees that computation time linearly increases with the number of template rotations and scalings tested for. This timings-experiment is performed on the pupil detection application, and the shown timings are only of *step 5* of the full detection pipeline (see Subsec. D.2.1 and Table 1) as this is the only step that is affected by the rotation/scale invariant extension.

TABLE 10
Average template matching results (± standard deviation, number of fails between parenthesis) for optic nerve head (ONH), fovea, and pupil detection in 5-fold cross validation.

Method	Success rate
ONH Detection	(1737 images)
No invariance	99.83% ± 0.26% (3)
Rotation invariance	99.60% ± 0.16% (7)
Fovea Detection	(1616 images)
No invariance	99.32% ± 0.26% (11)
Rotation invariance	97.10% ± 0.65% (47)
Pupil Detection	(1521 images)
No invariance	95.86% ± 0.89% (63)
Rotation invariance	94.48% ± 1.62% (84)
Scale invariance	95.33% ± 1.46% (71)
Rotation + scale invariance	94.28% ± 2.10% (87)

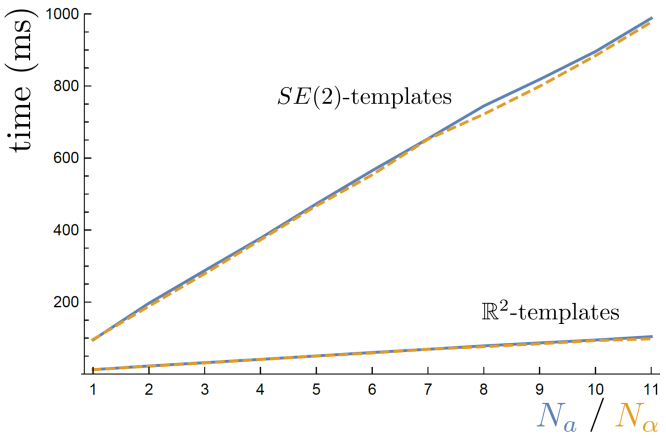

Fig. 12. Average computation times for the detection of one pupil/eye per image, using $SE(2)$ or $\mathbb{R}^2$ templates, testing for different template orientations or scalings. Two experiments are shown, in blue the number of template orientations $N_\alpha = 1$ and the number of scalings N_a is varied, in orange-dashed the number of scalings $N_a = 1$ and the number of rotations is N_α is varied.

REFERENCES

[1] J.-C. Yoo and T. Han, "Fast normalized cross-correlation," *CSSP*, vol. 28, no. 6, pp. 819–843, 2009.
[2] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in *CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, vol. 1, 2001, pp. I–511–I–518.
[3] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *CVPR 2005. IEEE Computer Society Conference on*, vol. 1, June 2005, pp. 886–893.
[4] D. Lowe, "Object recognition from local scale-invariant features," in *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, vol. 2, 1999, pp. 1150–1157 vol.2.
[5] H. Bay, T. Tuytelaars, and L. Van Gool, "Surf: Speeded up robust features," in *ECCV 2006*, ser. Lecture Notes in Computer Science, A. Leonardis, H. Bischof, and A. Pinz, Eds. Springer Berlin Heidelberg, 2006, vol. 3951, pp. 404–417.
[6] R. Duits, M. Felsberg, G. H. Granlund, and B. M. ter Haar Romeny, "Image analysis and reconstruction using a wavelet transform constructed from a reducible representation of the euclidean motion group," *IJCV*, vol. 72, no. 1, pp. 79–102, 2007.
[7] E. Bekkers, R. Duits, T. Berendschot, and B. ter Haar Romeny, "A multi-orientation analysis approach to retinal vessel tracking," *JMIV*, pp. 1–28, 2014.
[8] A. E. Hoerl and R. W. Kennard, "Ridge regression: Biased estimation for nonorthogonal problems," *Technometrics*, vol. 12, no. 1, pp. 55–67, 1970.

[9] J. Zhang, R. Duits, G. Sanguinetti, and B. M. ter Haar Romeny, "Numerical approaches for linear left-invariant diffusions on se(2), their comparison to exact solutions, and their applications in retinal imaging," *Numerical Mathematics: Theory, Methods and Applications*, vol. 9, pp. 1–50, 2 2016.

[10] R. Duits and E. Franken, "Left-invariant parabolic evolutions on SE(2) and contour enhancement via invertible orientation scores part I: Linear left-invariant diffusion equations on SE(2)," *Quarterly of Applied Mathematics*, vol. 68, no. 2, pp. 255–292, 2010.

[11] T. J. Hastie, R. J. Tibshirani, and J. H. Friedman, *The elements of statistical learning : data mining, inference, and prediction*, ser. Springer series in statistics. New York: Springer, 2009, autres impressions : 2011 (corr.), 2013 (7e corr.).

[12] M. Hebiri and S. van de Geer, "The smooth-lasso and other 1+ 2-penalized methods," *Electron J Stat*, vol. 5, pp. 1184–1226, 2011.

[13] R. Cuingnet, J. Glaunes, M. Chupin, H. Benali, and O. Colliot, "Spatial and anatomical regularization of svm: A general framework for neuroimaging data," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 35, no. 3, pp. 682–696, March 2013.

[14] H. Xu, C. Caramanis, and S. Mannor, "Robustness and regularization of support vector machines," *The Journal of Machine Learning Research*, vol. 10, pp. 1485–1510, 2009.

[15] A. A. Qazi, D. R. Jrgensen, M. Lillholm, M. Loog, M. Nielsen, and E. B. Dam, "A framework for optimizing measurement weight maps to minimize the required sample size," *Medical Image Analysis*, vol. 14, no. 3, pp. 255 – 264, 2010.

[16] C. Li and H. Li, "Network-constrained regularization and variable selection for analysis of genomic data," *Bioinformatics*, vol. 24, no. 9, pp. 1175–1182, 2008.

[17] C. De Boor, "A practical guide to splines," *Math. Comp.*, 1978.

[18] P. Green and B. Silverman, *Nonparametric regression and generalized linear models: a roughness penalty approach*. CRC Press, 1993.

[19] M. Unser, "Splines: a perfect fit for signal and image processing," *Signal Processing Magazine*, vol. 16, no. 6, pp. 22–38, Nov 1999.

[20] M. Unser, A. Aldroubi, M. Eden, and L. Fellow, "B-spline signal processing: Part i-theory," *IEEE TSP*, vol. 41, pp. 821–833, 1993.

[21] P. Thomas Fletcher, "Geodesic regression and the theory of least squares on riemannian manifolds," *International Journal of Computer Vision*, vol. 105, no. 2, pp. 171–185, 2013.

[22] N. Miolane and X. Pennec, "A survey of mathematical structures for extending 2D neurogeometry to 3D image processing," in *MICCAI-MVC*, Munich, Germany, Oct. 2015.

[23] X. Pennec, "Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements," *Journal of Mathematical Imaging and Vision*, vol. 25, no. 1, pp. 127–154, 2006.

[24] R. Vidal, Y. Ma, and S. Sastry, "Generalized principal component analysis (gpca)," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 27, no. 12, pp. 1945–1959, Dec 2005.

[25] O. Tuzel, F. Porikli, and P. Meer, "Learning on lie groups for invariant detection and tracking," in *CVPR*. IEEE, 2008, pp. 1–8.

[26] N. Patton, T. M. Aslam, T. MacGillivray, I. J. Deary, B. Dhillon, R. H. Eikelboom, K. Yogesan, and I. J. Constable, "Retinal image analysis: concepts, applications and potential," *Progress in retinal and eye research*, vol. 25, no. 1, pp. 99–127, 2006.

[27] S. A. Ramakanth and R. V. Babu, "Approximate nearest neighbour field based optic disk detection," *Computerized Medical Imaging and Graphics*, vol. 38, no. 1, pp. 49 – 56, 2014.

[28] M. E. Gegundez-Arias, D. Marin *et al.*, "Locating the fovea center position in digital fundus images using thresholding and feature extraction techniques," *Computerized Medical Imaging and Graphics*, vol. 37, no. 56, pp. 386 – 393, 2013, retinal Image Analysis.

[29] D. W. Hansen and Q. Ji, "In the eye of the beholder: A survey of models for eyes and gaze," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 32, no. 3, pp. 478–500, 2010.

[30] R. Duits, M. Duits, M. Almsick, and B. Haar Romeny, "Invertible orientation scores as an application of generalized wavelet theory," *PRIA*, vol. 17, no. 1, pp. 42–75, 2007.

[31] E. Bekkers, R. Duits, and B. ter Haar Romeny, "Optic nerve head detection via group correlations in multi-orientation transforms," in *Image Analysis and Recognition*, ser. LNCS, A. Campilho and M. Kamel, Eds. Springer, 2014, pp. 293–302.

[32] E. Bekkers, R. Duits, and M. Loog, "Training of templates for object recognition in invertible orientation scores: Application to optic nerve head detection in retinal images," in *EMMCVPR*. Springer, 2015, pp. 464–477.

[33] G. Citti and A. Sarti, "A cortical based model of perceptual completion in the roto-translation space," *JMIV*, vol. 24, no. 3, pp. 307–326, 2006.

[34] R. Duits and B. Burgeth, "Scale spaces on lie groups," in *Scale Space and Variational Methods in Computer Vision*, ser. Lecture Notes in Computer Science, F. Sgallari, A. Murli, and N. Paragios, Eds. Springer Berlin Heidelberg, 2007, vol. 4485, pp. 300–312.

[35] P. Craven and G. Wahba, "Smoothing noisy data with spline functions," *Num. Math.*, vol. 31, no. 4, pp. 377–403, 1978.

[36] F. O'sullivan, B. S. Yandell, and W. J. Raynor Jr, "Automatic smoothing of regression functions in generalized linear models," *JASA*, vol. 81, no. 393, pp. 96–103, 1986.

[37] C. Gu, "Cross-validating non-gaussian data," *Journal of Computational and Graphical Statistics*, vol. 1, no. 2, pp. 169–179, 1992.

[38] D. Xiang and G. Wahba, "A generalized approximate cross validation for smoothing splines with non-gaussian data," *Statistica Sinica*, vol. 6, pp. 675–692, 1996.

[39] S. Lu and J. Lim, "Automatic optic disc detection from retinal images by a line operator," *Biomedical Engineering, IEEE TBME*, vol. 58, no. 1, pp. 88–94, 2011.

[40] A. Aquino, M. E. Geg, and D. Mar, "Automated optic disk detection in retinal images of patients with diabetic retinopathy and risk of macular edema," *IJBLS*, vol. 8, no. 2, pp. 87–92, 2012.

[41] B. Dashtbozorg, A. M. Mendona, and A. Campilho, "Optic disc segmentation using the sliding band filter," *Computers in Biology and Medicine*, vol. 56, pp. 1 – 12, 2015.

[42] S. Sekhar, F. E. Abd El-Samie, P. Yu, W. Al-Nuaimy, and A. K. Nandi, "Automated localization of retinal features." *Applied Optics*, vol. 50, no. 19, pp. 3064–75, 2011.

[43] D. Marin, M. E. Gegundez-Arias *et al.*, "Obtaining optic disc center and pixel region by automatic thresholding methods on morphologically processed fundus images," *Computer Methods and Programs in Biomedicine*, vol. 118, no. 2, pp. 173 – 185, 2015.

[44] H. Yu *et al.*, "Fast localization and segmentation of optic disk in retinal images using directional matched filtering and level sets." *IEEE TITB*, vol. 16, no. 4, pp. 644–57, 2012.

[45] A. Giachetti, L. Ballerini, E. Trucco, and P. Wilson, "The use of radial symmetry to localize retinal landmarks," *Computerized Medical Imaging and Graphics*, vol. 37, no. 56, pp. 369 – 376, 2013, retinal Image Analysis.

[46] A. Youssif *et al.*, "Optic disc detection from normalized digital fundus images by means of a vessels' direction matched filter." *IEEE TMI*, vol. 27, no. 1, pp. 11–8, 2008.

[47] M. Foracchia *et al.*, "Luminosity and contrast normalization in retinal images." *MEDIA*, vol. 9, no. 3, pp. 179–90, 2005.

[48] J. Bruna and S. Mallat, "Invariant scattering convolution networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 8, pp. 1872–1886, 2013.

[49] W. Freeman and E. Adelson, "The design and use of steerable filters," *IEEE TPAMI*, vol. 13, no. 9, pp. 891–906, Sep 1991.

[50] A. Aquino, M. E. Gegundez-Arias, and D. Marin, "Detecting the optic disc boundary in digital fundus images using morphological, edge detection, and feature extraction techniques," *IEEE TMI*, vol. 29, no. 11, pp. 1860–9, 2010.

[51] S. Lu, "Accurate and efficient optic disc detection and segmentation by a circular transformation," *IEEE TMI*, vol. 30, no. 12, pp. 2126–2133, Dec 2011.

[52] M. Abramoff and M. Niemeijer, "Mass screening of diabetic retinopathy using automated methods," in *Teleophthalmology in Preventive Medicine*, G. Michelson, Ed. Springer, 2015, pp. 41–50.

[53] A. Aquino, "Establishing the macular grading grid by means of fovea centre detection using anatomical-based and visual-based features," *Comp. in Biology and Medicine*, vol. 55, pp. 61 – 73, 2014.

[54] H. Yu, S. Barriga, C. Agurto, S. Echegaray *et al.*, "Fast localization of optic disc and fovea in retinal images for eye disease screening," *Proc. SPIE*, vol. 7963, pp. 796 317–796 317–12, 2011.

[55] M. Niemeijer, M. D. Abràmoff, and B. van Ginneken, "Fast detection of the optic disc and fovea in color fundus photographs." *MEDIA*, vol. 13, no. 6, pp. 859–70, 2009.

[56] M. Leo, D. Cazzato, T. De Marco, and C. Distante, "Unsupervised eye pupil localization through differential geometry and local self-similarity matching," *PloS one*, vol. 9, no. 8, p. e102829, 2014.

[57] N. Markuš, M. Frljak, I. S. Pandžić, J. Ahlberg, and R. Forchheimer, "Eye pupil localization with an ensemble of randomized trees," *Pattern recognition*, vol. 47, no. 2, pp. 578–587, 2014.

[58] B. Kroon, A. Hanjalic, and S. M. Maas, "Eye localization for face matching: is it always useful and under what conditions?" in

Proceedings of the 2008 international conference on Content-based image and video retrieval. ACM, 2008, pp. 379–388.
[59] S. Asteriadis, N. Nikolaidis, and I. Pitas, "Facial feature detection using distance vector fields," *Pattern Recognition*, vol. 42, no. 7, pp. 1388–1398, 2009.
[60] P. Campadelli, R. Lanzarotti, and G. Lipori, "Precise eye and mouth localization," *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 23, no. 03, pp. 359–377, 2009.
[61] F. Timm and E. Barth, "Accurate eye centre localisation by means of gradients." in *VISAPP*. SciTePress, 2011, pp. 125–130.
[62] R. Valenti and T. Gevers, "Accurate eye center location through invariant isocentric patterns," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 34, no. 9, pp. 1785–1798, 2012.
[63] K. Yosida, "Resolvent and spectrum," in *Functional Analysis*, ser. Classics in Mathematics. Springer Berlin Heidelberg, 1995, vol. 123, pp. 209–231.
[64] L. Hörmander, "Hypoelliptic second order differential equations," *Acta Mathematica*, vol. 119, no. 1, pp. 147–171, 1967.
[65] E. P. Hsu, *Stochastic analysis on manifolds*. American Mathematical Soc., 2002, vol. 38.